Growth-at-risk and macroprudential policy design

Javier Suarez*
CEMFI, Casado del Alisal 5, 28014 Madrid, Spain. Email: suarez@cemfi.es

April 2022

Abstract

This paper explores the foundations for the application of the empirical growth-at-risk (GaR) approach to the assessment and design of macroprudential policies. It starts considering a stylized benchmark linear specification of the empirical GaR approach in combination with a linear-quadratic social welfare criterion that rewards expected GDP growth and penalizes the gap between expected GDP growth and GaR. If the growth rate follows a normal distribution, this welfare criterion can be microfounded as consistent with expected utility maximization under preferences for GDP levels exhibiting constant absolute risk aversion. The benchmark formulation implies an optimal policy rule linear in the risk indicator and an optimal gap between expected growth and GaR that does not depend on the time-varying risk indicator and is inversely related to the cost-effectiveness of macroprudential policy and the risk preference parameter. Extensions of the benchmark formulation show the potential to adapt the analysis and its insights to the richer specifications typically considered in empirical work.

JEL Classification: G01, G20, G28

Keywords: macroprudential policy, policy stance, growth-at-risk, quantile regressions

*I would like to thank Emilia Bonaccorsi di Patti, Stephen Cecchetti, Michal Dvorak, Stephan Fahr, Martin Gächter, Simone Manganelli, Tuomas Peltonen, Rafael Repullo, Abdelaziz Rouabah, Antonio Sanchez, Lars Svensson, two anonymous referees, participants at meetings of the Advisory Scientific Committee (ASC) and other structures of the European Systemic Risk Board (ESRB) as well as several seminars and conferences where prior versions of this work were presented for their helpful comments and suggestions. The views expressed herein are the exclusive responsibility of the author and do not necessarily reflect those of the institutions to which the author is affiliated. This work was supported by Spanish Government grant PGC2018-097044-B-I00. Declarations of interest: none.

1 Introduction

This paper is motivated by the growing attention paid to growth-at-risk (GaR) in the assessment of macroprudential policies. The concept arose as a natural extension of value-at-risk —a popular risk management concept— to the assessment of systemic risk. The value-at-risk of a given portfolio position is the critical level of the estimated distribution of possible losses over a reference horizon that realized losses will not exceed with a high probability (such as 95% or 99%) known as the confidence level of the assessment. Estimating the value-at-risk allows the portfolio holder to assess, for example, the capital position that would be needed to absorb the potential losses over the reference horizon with this confidence level of probability. From a statistical viewpoint, the value-at-risk of a portfolio is just the estimate of a low quantile (5% or 1% in the above examples) of the distribution of the value of the portfolio by the end of the reference horizon.

In parallel to the concept of value-at-risk, the GaR of an economy over a given horizon is a low quantile of the distribution of the (projected) GDP growth rate over the same horizon. In other words, the growth rate at which the probability of the realized growth rate falling below it equals a low benchmark level such as 10% or 5%.[1] In contrast to the standard macroeconomic focus on the expected value (and, perhaps, the variance) of aggregate output growth, looking at low quantiles of such growth implies, as in risk management, a focus on the severity of potential adverse outcomes. In addition to measuring this severity, the approach can provide information on the variables that determine the probability or severity of bad outcomes, including policy variables that might then be used to influence or "manage" the aggregate risk.

The rising popularity of GaR in financial stability and macroprudential policy assessments is driven by demand and supply factors. From the demand side, macroprudential policy assessment and design is in need of a quantitative framework that provides a baseline for policymakers' discussions, decision-making, and communication with the public similar to that provided by standard macroeconomic models, targets, and indicators in the field of monetary policy. For macroprudential policies, the multiplicity of tools, the multidimensional nature (and still vaguely defined concept) of systemic risk, data limitations, and the relatively short historical experience with the use of most policy tools pose significant challenges for the development of such a framework. As a result, macroprudential policy is largely assessed and developed following a piece-meal approach (that is, splitting the task by sector, tool, risk

[1] The use of lower implied confidence levels (90%, 95% in the examples above) in GaR than in value-at-risk is partly related to the fact that GDP is not observed at frequencies that allow for an accurate estimation of extremely low quantiles.

or detected vulnerability, or by using a combination of approaches) and relying on expert judgement for the qualitative integration of the pieces. While the aim is to cover the whole financial system, the resulting assessment is often less complete, integrated, systematic, and quantitative than in other policy fields.

From the supply side, the impulse to the concepts of GDP-at-risk (Cecchetti, 2008) and GaR (Adrian et al., 2018) came from their empirical suitability. It is related to the availability of econometric techniques that extend regression analysis (single dependent variable models, panel data models and vector auto-regressive models) to quantiles. Quantile regression techniques allow the focus to be shifted from modelling the conditional mean of the dependent variable to modeling the conditional quantiles, and potentially the whole conditional distribution, of the dependent variables.

These techniques allowed Cecchetti and Li (2008) to use the concept of GDP-at-risk as an empirically-viable summary measure of the impact of asset price booms on financial stability. This approach was further developed and promoted by the influential paper published by Adrian et al. (2019), which shows that the lower quantiles of the distribution of the US GDP growth rate fluctuate more and are more influenced by financial conditions than the upper quantiles, thus supporting the focus of macroprudential surveillance and policies on the lower quantiles. Adrian et al. (2018) documented the existence of a “term structure” in the effects of changes in financial conditions on GaR and interpreted their results as suggestive of potential policy trade-offs if macroprudential policies could counter the effects of changes in financial conditions.

Other contributions following a quantile-regression approach to the analysis of the impact of financial conditions and macroprudential policies on growth vulnerabilities include Caldera-Sánchez and Röhn (2016), De Nicolo and Lucchetta (2017), Prasad et al. (2019), Arbatli-Saxegaard et al. (2020), Chavleishvili et al. (2020), Brandao-Marques et al. (2020), Duprey and Ueberfeldt (2020), Franta and Gambacorta (2020), Figueres and Jarocinski (2020), Galán (in press), and Aikman et al. (2021).

Taken as a whole, establishing a stable body of evidence on the effects of macroprudential policies is still work in progress.[2] Part of this empirical work puts the emphasis on the capacity of financial variables to forecast low quantiles of GDP growth (and not necessarily high quantiles in a symmetric manner), thus suggesting a connection between financial factors or financial stability indicators and downside risk to output growth.[3] Other contributions

[2]For a survey covering earlier contributions and alternative approaches, see Galati and Moessner (2018).

[3]However, Plagborg-Møller et al. (2020) question the short-term forecasting capacity gained by considering variables such as the national financial conditions index (NFCI) in the prediction of GDP growth moments other than the conditional mean, while Brownlees and Souza (2021) challenge the out-of-sample

focus on the impact of macroprudential policies on GaR. For instance, Duprey and Ueberfeldt (2020), with data from Canada, find that the growth of credit to households contributes to tail risk and that the tightening of macroprudential policy (as captured by a qualitative index of policy actions) reduces tail risk but possibly at the cost of reducing mean GDP growth.[4] With a different approach, Gadea Rivas et al. (2020) find that credit growth increases the deepth of future recessions but also the duration of expansion periods thus suggesting the existence of a trade-off for the design of policies aimed to tame credit cycles.

Franta and Gambacorta (2020) also find positive financial stability implications of policy actions relating to the tightening of loan-to-value ratios and the provisioning of loan losses in a sample of 52 countries but they find no evidence of a cost in terms of mean growth outcomes. Likewise, Aikman et al. (2021) find that higher bank capitalization improves GaR over a three-year horizon without significantly reducing mean growth. In contrast, the results Galán (in press) are consistent with the view that the positive effect of macroprudential policy on tail outcomes might come at the expense of a negative effect of tightening actions on mean growth.

Brandao-Marques et al. (2020) go beyond the specific focus on mean or median growth and low-tail growth outcomes and assess the net benefits of macroprudential policies (as well as monetary policy and the combination of both policies) using a quadratic loss function of the type found in New Keynesian models: one that penalizes a weighted average of deviations of GDP growth with respect to potential growth and the square of the inflation rate. The authors conclude that, in the analyzed dataset, macroprudential policy is effective in the sense of producing, on average, a cumulative reduction in the expected value of the loss function (relative to a no-policy scenario) at all the horizons considered in the analysis.[5]

Conceptually, relative to other indicators of financial stability, GaR features the advantage of having an explicit and intuitive statistical interpretation and being measured in the same units as GDP growth, the most universal summary indicator of an economy's overall performance. Hence, quantitative contributions relating to the concept of GaR are followed with great interest (and some scepticism too) by the institutions involved in the assessment

short-term forecasting performance of quantile regressions relative to standard volatility models such as GARCH.

[4]The empirical analysis in Duprey and Ueberfeldt (2020) is complemented by a simple macroeconomic model that provides a microfoundation for the trade-off between mean growth and tail risk faced by macroprudential policy.

[5]To reach their conclusions, Brandao-Marques et al. (2020) address the endogeneity of historically-observed macroprudential policies (and issue that may limit the causal inference of several of the conmpeting papers) by relying on "policy shocks" defined as the unpredictable part of changes in the policy variables. Their analysis allows for but does not directly examine the existence or inexistence of a policy trade-off between mean and low-tail growth outcomes such as the one considered in the present paper.

and design of macroprudential policies. Policy-making institutions see in the GaR approach a promising avenue for the development of an integrated quantitative framework for macroprudential policy assessment and design.[6] However, as further discussed in Cecchetti and Suarez (2021), existing empirical efforts still lack a clear fit with an explicit policy design problem of the type considered for other macroeconomic policies (e.g., in the derivation of an optimal monetary policy rule).[7]

This paper addresses the fit of a benchmark GaR-type model into a canonical policy design problem. The aim is to provide a conceptual framework in which to fit the growing number of empirical contributions in the field. Starting with a stylized representation of the type of equations that the quantile regression approach may deliver (and considering afterwards multiple extensions of empirical relevance), the paper studies theoretically how macroprudential policy could be designed and evaluated within such an approach. Optimal policies are defined as those that maximize a social welfare criterion that rewards expected GDP growth and penalizes the gap between expected GDP growth and GaR and that, under specific conditions, can be microfounded as consistent with expected utility maximization under risk averse preferences for GDP levels. The paper characterizes the properties of optimal macroprudential policy rules in the benchmark setup and a number of extensions that bring the setup closer to what might be relevant in empirical applications and/or show the robustness and compatibility of the approach with alternative approaches. Implications are drawn on the possibility of assessing macroprudential policy stance with a metric emanating from the estimated equations of the empirical GaR approach.

The benchmark formulation abstracts from the time dimension by considering cumulative growth over the relevant policy horizon (e.g. two or three years) and focuses on the case in which macroprudential policy design is facing a trade-off: the available policy instrument can linearly increase GaR but at the expense of reducing expected GDP growth (e.g. because the tightening of some prudential requirement reduces growth vulnerabilities but has an expected contractive impact on economic activity).[8] We find that, under the benchmark formulation, the optimal policy rule is linear in a variable named the "risk indicator" which summarizes

[6] See, for instance, ESRB (2021).

[7] An exception in this landscape is Brandao-Marques et al. (2020) who, as previously mentioned, assess macroprudential policies against a loss function of the type used in New Keynesian models. However, their paper does not directly address the optimal policy design problem.

[8] The description of the macroprudential policy problem as one in which the policy maker faces a frontier in the mean growth vs. GaR (or tail risk) space can also be explicitly found in existing literature, including Aikman et al. (2018) and Duprey and Ueberfeldt (2020). However, these contributions do not elaborate on the social welfare criterion that is relevant in such a setting or on the properties of the implied optimal policies. Previously, Poloz (2014) referred in purely narrative/graphical terms to a policy frontier between financial stability risk and inflation-target risk.

the exogenous drivers of systemic risk. The sensitivity of the optimal policy to changes in this risk indicator turns out to be independent of the risk preferences embedded in the welfare criterion. This sensitivity depends directly on the impact of risk on the gap between expected growth and GaR, and inversely on the effectiveness of policy in reducing this gap. We also find that optimal macroprudential policy targets a gap between expected growth and GaR which does not depend on the level of the risk indicator but on the cost-effectiveness of macroprudential policy and the risk preference parameter.

The explored variations in the benchmark setup extend the analysis in several empirically or conceptually relevant dimensions. They cover cases (i) in which the policy variable has non-linear and/or state-contingent effects, (ii) with multiple policy variables and interactions between policy tools (including, possibly, monetary policy), and (iii) with discrete policy variables. Additional extensions in the discussion section consider (iv) the case of policies which seem to involve no trade-off between mean growth and GaR, (v) the compatibility of the GaR framework (and the main insights from the basic formulation) with the view that macroprudential policy involves various well-identified intermediate objectives each of which can be associated with one or a subset of targeted policy tools, (vi) the relationship with a framework focused on preventing and mitigating the effects of systemic financial crises, and (vii) the explicit consideration of intertemporal effects in the policy design problem. These extensions help analyze the robustness of the results of the benchmark formulation and to qualify the importance and limitations of the approach and its connection with alternative approaches.

The paper is structured as follows. Section 2 describes the benchmark model: a basic linear formulation of the type of equations estimated when using GaR approach. Section 3 explains the welfare criterion that is maximized to obtain the optimal policies under that formulation, derives and establishes the properties of the optimal macroprudential policy rule, and discusses how the estimates associated with the empirical counterpart of the model equations could help assess the stance of macroprudential policy. Section 4 develops several direct extensions of the benchmark setup, generalizing its results for a variety of empirically and policy relevant cases. Sections 5 contains further extensions and discussions that help put the explored GaR approach into context, connect it with alternative approaches, and interpret potentially puzzling empirical results. Section 6 concludes the paper. The Appendix contains the microfoundations of the GaR-based welfare criterion used in the design of the optimal policies, discusses the extent to which, when departing from normality, the focus on the conditional mean and the conditional low tail of the GDP growth distribution over a given horizon could have advantages over an alternative focus on only the conditional mean and conditional variance of the growth distribution, and shows the possibility of reformulating the analysis around the concept of growth-given-stress rather than GaR.

2 A benchmark formulation of the GaR approach

A quantile regression approach can deliver equations for arbitrary quantiles of GDP growth over relevant horizons. Let us consider a stylized representation of this approach that consists of two estimated equations: one for the mean (or perhaps the median) of the cumulative GDP growth over the policy horizon (e.g. one, two or three years), denoted by $\bar{y}$, and another for a relevant low quantile of the cumulative GDP growth over the same horizon, y_c.[9] The subscript c in y_c identifies the threshold probability (or *confidence level*) at which GaR is measured. By definition, y_c satisfies

$$\Pr(y \leq y_c) = c. \tag{1}$$

This means that the probability of experiencing growth rates that are lower than y_c over the relevant horizon is just c. The confidence level c can be thought to be 5% or 10% so that y_c reflects how bad growth may be under adverse circumstances typically (but not exclusively) associated with systemic financial distress.

As an initial benchmark, we consider the simple case in which the conditional forecast for *expected growth* $\bar{y}$ and *GaR* y_c over the policy horizon are determined by

$$\bar{y} = \alpha + \beta x + \gamma z, \tag{2}$$

and

$$y_c = \alpha_c + \beta_c x + \gamma_c z, \tag{3}$$

where x is a unidimensional *risk indicator* or exogenous driver of systemic risk (e.g., a driver of excessive credit growth or any other factor that potentially contributes to the accumulation of financial imbalances) and z is a unidimensional macroprudential *policy variable* (e.g., a bank capital-based measure such as the countercyclical capital buffer, CCyB, in Basel III). Let us further assume that the endogeneity of z has been treated well enough to allow for γ and γ_c to be interpreted as the causal impact of variations in z on expected growth and GaR, respectively.

[9] As mentioned in the introduction, this formulation abstracts from the exact shape of the path followed by GDP growth within the policy horizon, implicitly assuming that what matters for welfare is cumulative growth over the whole horizon. In contrast with policies focused on business cycle frequencies, these longer horizons are consistent with the assessment of the costs of systemic financial crises found in, for example, Laeven and Valencia (2018). Nevertheless, the approach could be extended to explicitly consider dynamics within the policy horizon and beyond, see Subsection 5.4.

To start with, we focus on the case in which the risk driver x has a negative impact on GaR and a less negative (or even positive) impact on expected growth, while policy variable z has a positive impact on GaR but a negative impact on expected GDP growth. That is, we assume

$$\beta_c < \min\{0, \beta\} \text{ and } \gamma < 0 < \gamma_c. \tag{4}$$

The assumed signs of γ and γ_c imply that the policy measured by z involves a trade-off. For example, if x describes a driver of excessive credit growth and z is the CCyB rate, a trade-off can arise because increasing the CCyB rate reduces the final systemic risk implied by, for instance, a credit boom (e.g., the probability and implications of an abrupt reversal) but, at the same time, has a contractive impact on aggregate demand and, hence, on the central outlook. Finally, we assume that the variation ranges of x and z together with the values of the intercepts α and α_c, guarantee $y_c < \bar{y}$ over the relevant range (otherwise, the linearity in (2) and (3) might lead to $\bar{y} < y_c$ which would not make sense for low values of c).

In this setup, the risk indicator x should be thought of as an exogenous driver of risk and not the final systemic risk faced by the economy. Systemic risk would be the result of the interaction of the risk driver x and policy z put in place to mitigate or counter its impact on tail outcomes. Thus, in the linear formulation above, systemic risk should be thought of as proportional to $\beta_c x + \gamma_c z$ rather than directly and solely x. Subsection 5.3 elaborates on how the use of GaR to gauge the implications of financial stability on growth outcomes could be related to more structural view in which financial instability (or systemic risk) explicitly refers to the probability and severity of systemic financial crises and macroprudential policy is explicitly directed to preventing and mitigating the effects of such crises.

The benchmark linear specification above implies that policy z monotonically affects y_c and $\bar{y}$. What really matters for the validity of the insights from the analysis below is that the stipulated effects are locally true over the relevant range of variation in z. Subsection 4.1 extends the analysis to the empirically relevant case in which the policy variable interacts with the risk variable, potentially making the sign of the impact of z on y_c to switch depending on the level of risk (capturing, for example, the risk-reducing effect of increasing a countercyclical buffer when vulnerabilities are building up and the value of releasing it during a crisis).

The policy trade-off implied by the assumption that $\gamma < 0 < \gamma_c$ is consistent with the empirical findings in Duprey and Ueberfeldt (2020), Galán (in press), and Gadea Rivas el al. (2020).[10] However, as discussed in the introduction, other authors have found no significantly

[10] Additionally, although with different methodologies, some quantitative assessments related to microprudential regulations, and capital requirements in particular, have found that the rise in capital requirements

negative effects of macroprudential policies on mean growth. In fact, it is perfectly conceptually plausible that, at low levels of activation, policies that enhance financial stability also increase expected growth. In such cases, from the perspective of the welfare metrics explored below, the activation of those policies up to the point a trade-off arises would involve a win-win situation and, hence, be not at all controversial. As further discussed in Subsection 5.1, it is only once (and if) a trade-off arises that the policy design problem becomes non-trivial, which explains our focus on the case a trade-off (locally) exists.

While the benchmark model contains just one risk variable and a policy variable, an advanced reader might easily extend many of the derivations and claims below to cases in which x and z are vectors of risk drivers and policy variables, respectively. As further discussed in Subsection 4.2, extensions with multiple policy variables and interactions between them could be applied to the analysis of optimal policy mixes, including mixes of macroprudential policies with other policies such as monetary policy.

3 Social preferences and the optimal policy rule

To further formalize the policy design problem associated with the empirical GaR formulation, suppose that the policy maker has preferences on growth outcomes that can be represented by the social welfare function

$$W = \bar{y} - \frac{1}{2} w (\bar{y} - y_c)^2, \tag{5}$$

where $w > 0$ measures the aversion for financial instability, which is here proxied by the magnitude of the quadratic deviations of GaR with respect to expected growth. This welfare criterion can be justified on various grounds.

First, as shown in Section A.1 of the Appendix, the form of W is consistent with the maximization of the expected utility of a representative risk-averse agent whose utility depends on GDP levels. Specifically, if the agent has preferences for GDP levels that exhibit constant absolute risk aversion (CARA) λ, then (5) provides an exact representation of these preferences under a value of w which is directly proportional to λ. When the growth rate is not normally distributed or preferences are not CARA, the extent to which (5) is a good approximation to the relevant expected utility is an empirical question.[11]

increase financial stability at a cost in terms of credit and output. In the early assessment of the MAG (2010) these costs were estimated to be modest and purely short-term. The structural assessment based on calibrated macroeconomic models in Mendicino et al. (2018) and Elenev at al. (2021), among others, imply more sizeable and long-lasting costs in terms of credit and output and the latent trade-offs imply the existence of interior welfare-maximizing levels of the capital requirements.

[11] A policy report by Cecchetti and Suarez (2021) explores the accuracy with which (5) approximates

Second, beyond the exact expected-utility microfoundations of the specific normal case, the welfare criterion in (5) can also be defended in heuristic or axiomatic terms as the representation of the preferences of a policy maker that faces a potential trade-off between improving mean outcomes and reducing the severity of very bad outcomes. An interesting feature of (5) from this perspective is that the dislike for "very bad outcomes" is proportional to the square of the distance between the very bad outcomes y_c and the mean outcomes $\bar{y}$, where the latter would play the role of a reference level (or status quo point) similar to those emphasized in some non-expected-utility formulations of agents' preferences for risk. Specifically, from the perspective of prospect theory (Kahneman and Tversky, 1979), the coefficient w in (5) could be interpreted as capturing loss aversion rather than risk aversion.[12]

Remaining flexible about the justification of the form of the welfare criterion W is important because in the specific case in which y is normally distributed, social preferences and the policy problem could have also been formulated in the usual mean-variance terms of portfolio theory, with an equation describing the dependence of the standard deviation of the growth rate σ_y on x and z replacing (3) (see Section A.2 of the Appendix for details). What this means is that the true advantages of adopting a GaR approach (instead of a mean-variance approach) in the formulation of the macroprudential policy problem must derive from the fact that in reality: (i) the conditional distribution of y is not Gaussian, and (ii) as documented in recent empirical work, the financial factors and policy tools on which macroprudential policy focuses affect the conditional low quantiles of the true growth distribution in a stronger and more clearly identifiable manner than its conditional variance. From this perspective, an advantage of the quantile-regression approach to the modelling of the quantile y_c is that it does not require the assumption of a specific conditional distribution for the growth rate. In other words, nothing prevents the estimated version of equations (2) and (3) from capturing features such as the potential left skewness of the true conditional distribution of the GDP growth rate or effects of a policy variable on tail growth that would not be equivalently captured by modelling the effect of such variable on the variance of the growth rate.

the expected utility in a number of empirically motivated examples that depart from the normality and CARA assumptions. For realistic levels of variability of cumulative GDP growth over three-year periods, the accuracy of the metric provided by (5) is very good even if the distribution of the growth rate exhibits kurtosis and skewness and preferences feature constant relative risk aversion (CRRA) as commonly assumed in the macroeconomic literature.

[12] The asymmetric focus on low tail losses can also be related to Fishburn (1977) who explores preferences in which the decision maker is averse to obtaining below-target payoffs. Kilian and Manganelli (2008) analyze the decision problem of a central banker using that approach. In a similar vein, Svensson (2003) considers a monetary policy problem under preferences that asymmetrically penalize extreme events.

In any case, the macroprudential policy design implications obtained under (5) can be understood as a conceptual benchmark for more complicated or general formulations that could only be analyzed numerically. For example, the policy problem could be directly formulated as one of maximizing the expected value of the utility derived from growth outcomes taking the whole conditional distribution of those outcomes into account (that is, all the quantiles of the growth distribution). However, the characterization of optimal policies under that approach would have to be necessarily numerical.[13]

3.1 The optimal policy rule

Under our formulation, a macroprudential policy that is optimal conditional on a risk level x would maximize W given x. That is, it would be characterized by the policy rule

$$z(x) = \arg\max_{z} W(x, z), \tag{6}$$

where $W(x, z)$ describes W as a function of the risk indicator x and the policy variable z after taking into account the dependence of $\bar{y}$ and y_c on both variables as described in (2) and (3).

If the optimal policy is interior, it must solve the following first order condition (FOC):

$$\frac{\partial \bar{y}}{\partial z} - w(\bar{y} - y_c)\left(\frac{\partial \bar{y}}{\partial z} - \frac{\partial y_c}{\partial z}\right) = 0, \tag{7}$$

which uses the chain rule in (5). Under the benchmark specification, this FOC can be written as

$$\gamma - w(\alpha + \beta x + \gamma z - \alpha_c - \beta_c x - \gamma_c z)(\gamma - \gamma_c) = 0. \tag{8}$$

Solving for z leads to the *macroprudential policy rule*

$$z(x) = \phi_0 + \phi_1 x, \tag{9}$$

with

$$\phi_0 = \frac{\alpha - \alpha_c}{\gamma_c - \gamma} + \frac{\gamma}{w(\gamma_c - \gamma)^2} \tag{10}$$

and

$$\phi_1 = \frac{\beta - \beta_c}{\gamma_c - \gamma}. \tag{11}$$

Under the assumptions in (4), the *intercept* of the policy rule ϕ_0 can in principle have any sign since it is the sum of a first term which will most typically be positive (specifically if

[13] This is the approach adopted by Brandao-Marques et al. (2020) for the valuation of the loss function against which they assess the average net gains of historically-observed macroprudential policies.

$\alpha - \alpha_c > 0$) and a second term which is negative (since $\gamma < 0$). However, ϕ_0 is intuitively increasing in the policy maker's preference for financial stability w (since the absolute size of the negative term declines with w) and also increasing in the difference $\gamma_c - \gamma > 0$, which measures the *effectiveness* of the policy variable in reducing the *gap* between expected growth and GaR, $\bar{y} - y_c$.[14]

Interestingly, the parameter ϕ_1 which measures the *responsiveness* of the optimal policy to variations in risk indicator x is positive and independent of the preference parameter w. So in this setup, policy makers with different preferences for financial stability would differ in the level at which they use macroprudential policy but not in the extent to which they modify their policies in response to changes in the risk assessment. This optimal policy responsiveness is directly proportional to the *impact of risk* x on the gap between expected growth and GaR ($\beta - \beta_c$, which is positive under (4)) and inversely proportional to the *effectiveness of policy* z in reducing the gap ($\gamma_c - \gamma$, which is also positive under (4)).

Thus, while estimating equations (2) and (3) does not, per se, allow the preference parameter w to be estimated, it can provide an estimate of the optimal policy responsiveness parameter ϕ_1 and its components $\beta - \beta_c$ and $\gamma_c - \gamma$. It can also allow the optimal policy rule to be represented for different illustrative values of the preference parameter w.[15]

3.2 Graphical illustration of the policy problem

Further understanding of the interaction between the policy trade-offs implied by (2) and (3) and the preferences reflected in (5) can be obtained by depicting the frontier of pairs of y_c and $\bar{y}$ that can be reached, for a given value of the risk variable x, by varying the policy variable z. Mathematically this *conditional policy frontier* (for a given x) is defined by the line:

$$\bar{y} = \left(\alpha - \frac{\gamma}{\gamma_c}\alpha_c\right) + \left(\beta - \frac{\gamma}{\gamma_c}\beta_c\right) x + \frac{\gamma}{\gamma_c} y_c, \tag{12}$$

[14] For instance, in the polar case in which the policy maker has absolute preference for financial stability ($w \to \infty$), the intercept would become just $(\alpha - \alpha_c)/(\gamma_c - \gamma)$ and lead to a solution with $\bar{y} = y_c$ (which, although unrealistic in practice, is mathematically feasible given the linearity of (3) and (2) in x and z). In the other polar scenario with no preference for financial stability ($w \to 0$), we would have $\phi_0 \to -\infty$, implying that the policy maker would choose the lowest possible value of z, since under the linear specification of (2) and $\gamma < 0$ this is the way to maximize expected growth (albeit at the cost of minimizing GaR).

[15] Saying that the approach "can" deliver an estimate of ϕ_1 and the policy rule for illustrative values of w is a reminder of the importance of relying on estimates of parameters γ and γ_c that reflect the causal impact of policy on growth outcomes and not just some partial correlations between historical realizations of policy and outcomes.

which is downward sloping in $(y_c, \bar{y})$ space. Figure 1 depicts the policy frontier for a given value of x. The point $(y_c(x,0), \bar{y}(x,0))$ corresponds to the case in which $z = 0$.[16] Intuitively, with $\gamma_c > 0$, choosing $z > 0$ allows higher values of y_c to be obtained but, with $\gamma < 0$, this comes at the cost of lowering $\bar{y}$.

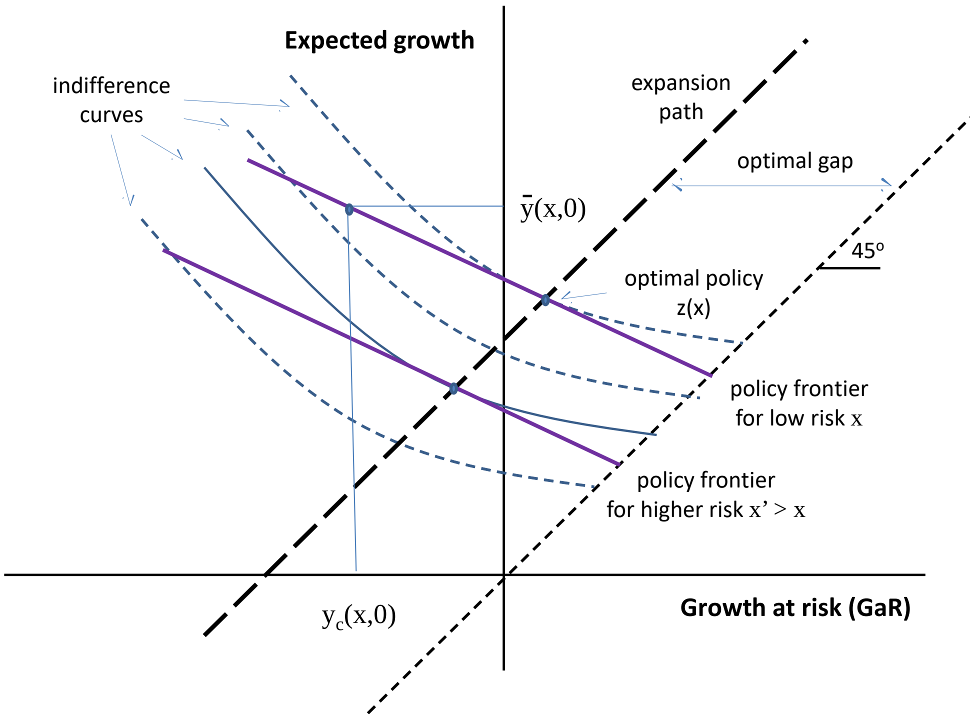

Figure 1. Graphical illustration of the policy problem

[NOTE TO TYPESETTERS: Figure intended for insertion as a two-column fitting image. JPG file with the graphical contents has been provided as a separate file.]

The preferences in (5) describe a map of indifference curves in $(y_c, \bar{y})$ space that are convex parabolas which reach their minima on the ray $y_c = \bar{y}$. The map makes economic sense to the left of this ray. Intuitively, for $w > 0$, on each indifference curve, any decline in y_c should be compensated with an increase in $\bar{y}$ in order to keep the welfare level unchanged. Moreover, for a given decline in y_c, the required compensating increase in $\bar{y}$ increases with

[16] Under the assumed linearity, there is nothing special about $z = 0$ but in specific applications the policy variable could be normalized so that it means something, e.g., the historical mean or "normal stance" of the corresponding policy (then $z < 0$ would represent a stance that is looser than normal and $z > 0$ a stance tighter than normal). For some policy instruments there may be a natural lower bound to z, e.g. a CCyB rate under Basel III cannot be negative (although in practice there are instances of capital forbearance that might be similar to having $z < 0$ for this instrument). Explicit consideration of these bounds would raise complications regarding occasionally binding constraints that are familiar in other contexts.

the distance from the ray $\bar{y} = y_c$. This explains why the FOC (7) includes the term $w(\bar{y} - y_c)$, which accounts for the marginal cost of financial instability.

Optimal policy $z(x)$ is the choice of z that leads to maximum welfare on the corresponding conditional policy frontier. In other words, $z(x)$ is the policy level that leads to the point where the conditional policy frontier is tangent to the map of indifference curves. The determinants of the slopes of these curves imply that, all other things being equal, a policy maker with a stronger preference for financial stability will choose, within the policy frontier, combinations of $(y_c, \bar{y})$ that involve a lower $\bar{y}$ and a higher y_c, that is, a lower gap between expected growth and GaR.

3.3 Optimal target gap property

What happens when the risk indicator x varies? From (12), changes in the risk indicator x cause a parallel shift in the policy frontier (just another implication of the linear formulation). In a scenario with $\beta > \gamma\beta_c/\gamma_c$, which might correspond to a situation of financial exuberance, the factors that drive risk up have such a strong effect on expected growth that the policy frontier shifts upwards, permitting better expected growth and GaR outcomes to be reached. In contrast, when $\beta < \gamma\beta_c/\gamma_c$, which might correspond to a situation of materialization of financial vulnerabilities, an increase in risk makes the frontier to shift downwards, depressing the combinations of mean and low-tail growth outcomes that can be reached.

In both cases, however, the optimal policy rule (9) of the benchmark model implies that the rise in x should be responded with a tightening of policy decision z, indicating that the deterioration of GaR y_c that would occur if policy were not adjusted is, at least partly, offset by increasing z. When risk has a positive marginal impact on expected growth ($\beta > 0$), the optimal policy response will diminish the raw positive effect of risk indicator x on mean growth $\bar{y}$. When risk has a negative marginal impact on expected growth ($\beta < 0$), the optimal policy response will still aim to offset its even more negative effect on y_c by lowering mean growth $\bar{y}$ beyond the impact implied by the raw negative effect of risk indicator x on $\bar{y}$.

Mathematically, the final impact of changes in risk indicator x on $\bar{y}$ and y_c can be seen by substituting the optimal policy rule $z(x)$ in (2) and (3), which leads to

$$\bar{y} = (\alpha + \gamma\phi_0) + (\beta + \gamma\phi_1)x = (\alpha + \beta\phi_0) + \frac{\gamma_c\beta - \gamma\beta_c}{\gamma_c - \gamma}x, \tag{13}$$

and

$$y_c = (\alpha_c + \gamma_c\phi_0) + (\beta_c + \gamma_c\phi_1)x = (\alpha_c + \beta_c\phi_0) + \frac{\gamma_c\beta - \gamma\beta_c}{\gamma_c - \gamma}x. \tag{14}$$

Interestingly, the coefficient of risk indicator x is identical in these two equations, which implies that the optimal policy rule would keep the gap between expected growth and GaR constant:

$$\bar{y} - y_c = (\alpha - \alpha_c) + (\gamma - \gamma_c)\,\phi_0 = \frac{1}{w}\frac{(-\gamma)}{\gamma_c - \gamma} = \frac{1}{w}\frac{1}{1 + \gamma_c/(-\gamma)} \tag{15}$$

Notice that this *target gap* is positive under (4) since $\gamma < 0$. The target gap decreases in the preference for financial stability w and increases in the *marginal growth-gap rate of transformation* implied by the policy frontier, $(-\gamma)/(\gamma_c - \gamma)$, which can be rewritten as $1/[1 + \gamma_c/(-\gamma)]$ to better visualize its negative dependence with respect to the *marginal cost-effectiveness* of the policy variable: the ratio of the c-quantile-improving effect $\gamma_c > 0$ to the mean-reducing effect $-\gamma < \gamma_c$.

In fact, this "constant target gap" property can be directly obtained from the FOC in (7), which can be rearranged as

$$\bar{y} - y_c = -\frac{1}{w}\frac{\frac{\partial \bar{y}}{\partial z}}{\frac{\partial \bar{y}}{\partial z} - \frac{\partial y_c}{\partial z}} = \frac{1}{w}\frac{1}{1 + \gamma_c/(-\gamma)}. \tag{16}$$

Graphically this implies that changes in risk indicator x and the optimal policy response under $z(x)$ describe a linear *expansion path* in $(y_c, \bar{y})$ with a slope equal to one. Thus, starting from the optimal policy identified in Figure 1 for a particular risk level x, changes in x will lead to combinations $(y_c, \bar{y})$ on the line with slope one that goes through that point.

More specifically, when risk does not increase expected growth too much (that is, in the case $\beta < \gamma\beta_c/\gamma_c$ described above), the coefficient of x in the reduced-form equations (14) and (13) is negative. Thus, when the risk indicator increases and policy responds optimally, expected growth and GaR *deteriorate* by the same amount, keeping the gap between expected growth and GaR, $\bar{y} - y_c$, constant. This is the case depicted in Figure 1 where the policy frontier under $x' > x$ lies on the left of the policy frontier for x.

Otherwise (when $\beta > \gamma\beta_c/\gamma_c$), the coefficient of x in (13) and (14) is positive, and rises in x and the optimal policy response lead expected growth and GaR to *improve* by the same amount but once again keeping $\bar{y} - y_c$ constant.

An important corollary to these findings is that, under the specified preferences, macroprudential policy should not target a constant GaR or keep GaR above a certain lower bound but should allow the GaR target to comove (actually by the same amount!) with the expected growth estimate. In other words, these derivations suggest that the $\bar{y} - y_c$ gap is a more useful indicator of macroprudential policy stance than each of its components separately.[17]

[17] If the GDP growth rate is normally distributed, the formulation of the macroprudential policy problem

3.4 A framework for policy assessment and design?

A tentative list of policy-relevant outputs that the framework developed so far can deliver is set out below.

1. Estimating (2) and (3) allows us to positively describe the direct impact of risk x and policy z on expected growth and GaR, as well as the policy trade-offs involved.

2. The policy trade-offs can be further illustrated using a policy frontier as shown in Figure 1.

   (a) If it is evaluated at the historical mean value of x, this frontier could be called the *mean policy frontier.*

   (b) Under the linear specification, the conditional policy frontier is just a parallel shift of the mean policy frontier. The relative position of the conditional frontier relative to the mean frontier may indicate whether the economy is facing a state of above-normal or below-normal risk exposure.

3. If the preferences of the policy maker can be described with a welfare criterion as in (5), the following additional points can be made:

   (a) The optimal policy responsiveness to the risk indicator can be measured using $\phi_1 = \frac{\beta - \beta_c}{\gamma_c - \gamma}$, as in (11). This coefficient is independent of parameter w that describes the policy maker's preference for financial stability.

   (b) If the policy maker follows the optimal policy rule, it will implicitly target a constant gap between expected growth and GaR, as in (16). From this gap, and the estimates of the empirical GaR model, the implicit preference parameter could be recovered (inferred) from the condition

$$w = \frac{1}{\bar{y} - y_c} \frac{1}{1 + \gamma_c/(-\gamma)}. \tag{17}$$

on the basis of growth-given-stress (that is, the expected growth rate conditional on the growth rate being below its c-quantile) would make no difference. This happens because in a normal distribution the distance between the mean and the c-quantile is proportional to the distance between the mean and the expected value of the random variable conditional on being below the c-quantile. This allows us to re-express the welfare criterion in (5) in terms of growth-given-stress while retaining the microfoundation of W (see section A.3 of the Appendix). Consequently, the constant target gap property of the optimal policy could also be expressed in terms of the distance between expected growth and growth-under-stress.

(c) Conditional on a reference value of the preference parameter w, the optimal policy rule can be fully described using (9). Graphically, it can be described using the expansion path previously illustrated in Figure 1.

(d) For each assessment of risk x, the optimal policy choice can be graphically described as the point of the conditional policy frontier that intersects with the expansion path.

(e) Conditional on an assessment of risk x, a policy stance could be deemed *inefficient* if it leads to points sufficiently far away from the policy frontier. However, when the policy variable x is unidimensional, all choices of x are "efficient," so the concept of inefficiency is only useful when there are two or more policy variables (as in some of the extensions discussed below).

(f) Conditional on the reference value of the preference parameter w and an assessment of risk x, a policy stance could be deemed *suboptimal* if it is sufficiently far away from the expansion path. This corresponds to an excessive distance between z and $z(x)$ or, in terms of outcomes, a $\bar{y} - y_c$ gap that is excessively far from its target. Thus policy would be *too tight* if z is sufficiently higher than $z(x)$ and, equivalently, if the gap $\bar{y} - y_c$ is well below the target. Conversely, policy would be *too loose* if z is sufficiently lower than $z(x)$ and, equivalently, if the gap $\bar{y} - y_c$ is well above the target.

In a multi-country environment, the empirical framework considered in this paper may involve country-specific versions of equations (2) and (3) as well as cross-country differences in the risk preference parameter w.[18] In the context of the "single country" baseline specification explored in this paper, these country differences can be thought of as just having different values of the involved parameters and their implications can be easily extracted from the prior discussion of the benchmark case.

In the next sections of the paper, we consider generalizations and variations of the benchmark model that bring the formulation closer to the one that might be relevant from an empirical perspective. The reader will notice that many of the concepts introduced for the benchmark case can be suitably adapted to each specific extension or variation of the setup.

[18] Obtaining the former does not necessarily mean running quantile regressions country by country. Alternatives include running panel quantile regressions that allow for country fixed effects or coefficients for the risk indicators or the policy variables that vary with some country-specific characteristics (e.g., obtained by interacting x and z with variables representing differences in the structure of countries' financial or legal systems).

In this sense, the analysis of the benchmark model, in spite of its simplicity, remain informative about how to address the assessment and design of macroprudential policy using an empirical GaR approach.

4 Generalizing the benchmark formulation

This section demonstrates the capacity of the GaR-based framework to accommodate multiple considerations of empirical and policy relevance, discussing the robustness of the properties of optimal macroprudential policies in each of the extensions.

4.1 Policy variable with non-linear and/or state-contingent effects

Extending the benchmark model to include non-linear effects of the policy variable or effects that vary with the state of the economy is of great practical importance. Relevant non-linearities may be related to the diminishing effectiveness of the policy variable (or variables, if there are several) in improving GaR or to a marginally increasing negative impact of the policy variable on expected growth. Both cases would intuitively lead to a gradual reduction in gains from policy tightening (increasing z) as the primary risk indicator x increases and to a target gap that, instead of invariant to x, increases with x. Yet another interesting case is when the (primary) risk variable x (or a secondary risk variable that captures the materialization of systemic risk) interacts with the policy variable reducing its effectiveness or even switching the direction in which policy might best respond to the situation (e.g. by releasing rather than replenishing a countercyclical capital buffer or loosening rather than tightening a borrower-based measure).

To avoid introducing too many complications at once, consider a generalized version of (2) and (3) in which risk and policy can still be represented by the unidimensional variables x and z but their impact on expected growth and GaR is of the general form

$$\bar{y} = Y(x, z), \tag{18}$$

and

$$y_c = Y^c(x, z), \tag{19}$$

where the functions $Y(\cdot)$ and $Y^c(\cdot)$ have first partial derivatives satisfying $Y_x^c < \min\{0, Y_x\}$ and $Y_z < 0 < Y_z^c$, and second partial derivatives satisfying $Y_{zz}^c < Y_{zz} < 0$. In this case, an interior optimal policy would solve the following FOC:

$$Y_z(x, z) - w[Y(x, z) - Y^c(x, z)][Y_z(x, z) - Y_z^c(x, z)] = 0, \tag{20}$$

where the dependence of $\bar{y} = Y(x, z)$ and $y_c = Y^c(x, z)$ on x and z has been made explicit to emphasize the type of non-linear equation that would have to be solved to find the optimal policy rule $z(x)$. It is immediate to check that having $Y^c_{zz} < Y_{zz} < 0$ (that is, a diminishing marginal impact of policy z on the gap) is sufficient to satisfy the relevant second order condition for a maximum.

By rearranging (20), we can obtain an expression for the gap associated with the optimal policy that is qualitatively very similar to that obtained for the benchmark model:

$$Y(x, z) - Y^c(x, z) = \frac{1}{w} \frac{1}{1 + Y^c_z(x, z)/(-Y_z(x, z))}. \tag{21}$$

Specifically, it leads to similar conclusions regarding how the social aversion to financial instability w and the marginal cost-effectiveness of the policy (here $Y^c_z(x, z)/(-Y_z(x, z))$) affect the optimal gap. However, in this case the target gap would not generally be invariant to the risk indicator x and the optimal policy may no longer respond monotonically to changes in x.

To see the latter, we can use the Implicit Function Theorem on (20). The full differentiation of the FOC with respect to x and z, together with the fact that the second order condition for a maximum holds, makes the sign of $z'(x)$ the same as the sign of the derivative of the left hand side of (20) with respect to x. Omitting the arguments of the functions $Y(x, z)$ and $Y^c(x, z)$ and their derivatives, the relevant expression is

$$S = Y_{zx} - w(Y_x - Y^c_x)(Y_z - Y^c_z) - w(Y - Y^c)(Y_{zx} - Y^c_{zx}), \tag{22}$$

whose sign depends, in general, on the cross derivatives Y_{zx} and Y^c_{zx}. To gain further insights on the possibilities emerging under this generalization, we next consider two illustrative cases.

4.1.1 When risk per se does not affect policy effectiveness

Suppose $Y_{zx} = Y^c_{zx} = 0$ so that, as in the benchmark model, risk per se does not affect the marginal effects of policy variable z on expected growth and GaR. In this case,

$$S = -w(Y_x - Y^c_x)(Y_z - Y^c_z) > 0, \tag{23}$$

so $z'(x) > 0$ as in the benchmark model. However, having $Y^c_{zz} < Y_{zz} < 0$ means that the marginal cost-effectiveness of the policy, $Y^c_z/(-Y_z)$, declines with its intensity z. Using (21), this implies that the target gap $Y - Y^c$ increases with x. In other words, as risk deteriorates, the policy maker would accommodate the diminishing cost-effectiveness of the policy tool by widening the targeted gap between expected growth and GaR.

4.1.2 When risk per se affects policy effectiveness

To explore a simple case in which the effectiveness of macroprudential variable is state contingent, consider an extension of the benchmark model in which expected growth is still determined by (2) but the GaR equation includes a new term with the interaction (product) of x and z:

$$y_c = \alpha_c + \beta_c x + \gamma_c z + \delta_c xz, \tag{24}$$

with $\delta_c < 0$. The FOC for an interior solution to the policy design problem would now be

$$\gamma - w(\alpha + \beta x + \gamma z - \alpha_c - \beta_c x - \gamma_c z - \delta_c xz)(\gamma - \gamma_c - \delta_c x) = 0, \tag{25}$$

which extends (8) with two new terms in δ_c and implies that the target gap would now be decreasing in x:

$$\bar{y} - y_c = \frac{1}{w}\frac{1}{1 + (\gamma_c + \delta_c x)/(-\gamma)}, \tag{26}$$

since, intuitively, $\delta_c < 0$ makes the marginal cost-effectiveness of the policy decline with x.

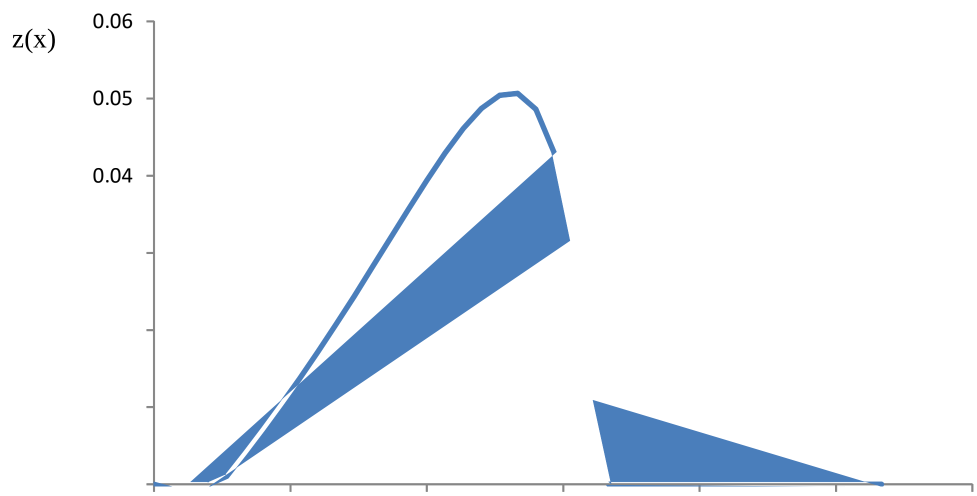

Figure 2. Optimal policy when risk diminishes policy effectiveness
This figure represents optimal policy $z(x)$ as a function of risk variable x in a specification based on equations (2) and (24). Parameter values: $\alpha = -\alpha_c = 0.2$, $\beta = 0.1$, $\beta_c = -0.5$, $\gamma = -0.2$, $\gamma_c = 2$, $\delta_c = -5$, and $w = 1.4784$.

[NOTE TO TYPESETTERS: Figure intended for insertion as a two-column fitting image. JPG file with the graphical contents has been provided as a separate file. The caption under the title of the figure reads as follows:

This figure represents optimal policy $z(x)$ as a function of risk variable x in a specification based on equations (2) and (24). Parameter values: $\alpha = -\alpha_c = 0.2$, $\beta = 0.1$, $\beta_c = -0.5$, $\gamma = -0.2$, $\gamma_c = 2$, $\delta_c = -5$, and $w = 1.4784$.]

In fact, solving for z leads to a modified version of the benchmark policy rule:

$$z(x) = \left(\frac{\alpha - \alpha_c}{\gamma_c + \delta_c x - \gamma} + \frac{\gamma}{w(\gamma_c + \delta_c x - \gamma)^2} \right) + \frac{\beta - \beta_c}{\gamma_c + \delta_c x - \gamma} x, \tag{27}$$

which is no longer linear (or even monotonically increasing) in x. The presence of the term in $\delta_c x$ in the second parenthesis in (25) means that there is a critical level of risk $\hat{x} = (\gamma_c - \gamma)/(-\delta_c) > 0$ above which increasing z would increase rather then decrease the gap between expected growth and GaR. So an optimal policy will stop increasing z with x before reaching that point. Figure 2 illustrates the non-monotonic policy rule for a numerical example in which the policy variable is constrained to take positive values.[19] It shows a case in which beyond a critical level of risk (of about 0.27 in the example) the optimally policy is drastically loosened as risk further increases.[20]

4.2 Interactions between multiple policy tools

The benchmark model considers just one policy variable for simplicity but the GaR-based framework can be extended to consider a vector of $j = 1, 2, ...M$ policy variables z_j, thus allowing to consider the case in which macroprudential policy can play with several tools as well as the interactions between macroprudential policy and other policies (e.g. monetary policy) that might represented by some of the M tools.

4.2.1 Linearity implies the existence of a dominant policy tool

Let us consider an extended version of (2) and (3) in which M different continuous policy variables z_j linearly affect y_c and $\bar{y}$ with coefficients γ_{cj} and γ_j, respectively. Assume that these coefficients satisfy $\gamma_j < 0 < \gamma_{cj}$ as in (4) and further assume that the variables are

[19] Linear-quadratic specifications of the equations for $\bar{y}$ and y_c, which have the advantage of being potentially estimated with standard linear quantile regression techniques, can be thought of as second-order local approximations to more general non-linear versions of $Y(x, z)$ and $Y^c(x, y)$. However, as it is well-known, they must be managed with care since for values of x and z away from the ranges where the approximation is valid, the quadratic terms can lead to non-sensical results. For instance, under the parameters behind Figure 2, for $x > \hat{x}$, the equations would imply that a large negative z would allow to reach a large value of $\bar{y}$ and, at the same time, an even larger value of y_c, which is inconsistent with y_c representing a low quantile.

[20] In the logic of this example, it is a sufficiently high value of x itself what makes policy tightening no longer an optimal response to higher values of x. Empirical applications might instead have the term in δ_c be made of the interaction of z (or xz) with a second risk indicator (even a binary variable) that signals the materialization of systemic risk.

scaled so that $z_j = 0$ is the relevant lower bound applicable to all of them. In this linear world, as exploring the relevant FOCs reveals, there will generally be one variable dominating the others in the maximization of W. This most efficient or *preferred policy tool* j^* would be the one featuring the lowest value of what was referred to as the marginal growth-gap rate of transformation in the single policy variable benchmark,

$$\frac{\frac{\partial \bar{y}}{\partial z_j}}{\frac{\partial \bar{y}}{\partial z_j} - \frac{\partial y_c}{\partial z_j}} = \frac{1}{1 + \gamma_{cj}/(-\gamma_j)} > 0, \tag{28}$$

that is, the policy tool with the best marginal cost-effectiveness as measured by $\gamma_{cj}/(-\gamma_j)$.[21] Intuitively, when $\gamma_{cj}/(-\gamma_j)$ is higher the same reduction in the gap between expected growth and GaR can be achieved at a lower cost in terms of expected growth. For the most efficient tool, the optimal value of z_{j^*} would be the value that satisfies the counterpart of equation (7). The associated policy rule would be the same as in (9) with ϕ_0 and ϕ_1 particularized to the preferred tool j^*. All elements in Figure 1 remain valid if the policy frontiers are also particularized to those obtained using the dominant tool.

4.2.2 Non-linearities can give rise to optimal policy mixes

The optimality of using non-trivial combinations of tools in macroprudential policy would only emerge under departures from linearity.[22] For example, optimal policies could be obtained that involve using several tools at the same time if the effectiveness of each policy tool in reducing GaR is marginally decreasing (that is, their impact on GaR is, for instance, given by some functions $\Gamma_{cj}(z_j)$ with $\Gamma'_{cj} > 0$ and $\Gamma''_{cj} < 0$) or if there are complementarities between tools under a general quasi-concave function $\Gamma_c(z_1, z_2, ...z_M)$ that replaces the terms $\Sigma_{j=1}^{M} \gamma_{cj} z_j$ in the extended version of (3).

In such a non-linear world, all policy variables activated at a strictly positive level at the optimum would satisfy a properly modified version of (7) and, consequently, (16) implying

$$\bar{y} - y_c = -\frac{1}{w} \frac{\frac{\partial \bar{y}}{\partial z_j}}{\frac{\partial \bar{y}}{\partial z_j} - \frac{\partial y_c}{\partial z_j}} = \frac{1}{w} \frac{1}{1 + \frac{\partial \Gamma_c}{\partial z_j}/(-\gamma_j)}. \tag{29}$$

[21] All the other policy variables should remain at their lower bound. In terms of Figure 1, using an inferior tool would imply moving over policy "frontiers" that also go through the point $(y_c(x,0), \bar{y}(x,0))$ but with steeper slopes, confirming that such tool would only be able to increase y_c by causing larger declines in $\bar{y}$. Conditional on using a less cost-effective tool, equation (16) would imply that the target gap should be larger, thus accommodating the harder trade-off faced along the corresponding policy frontier.

[22] As illustrated in the discussion of the case with interactions between the effects of x and z, non-linearities can be captured in a linear quantile regression setup by non-linearly transforming the variables entering the right hand side of the regression (e.g. by having quadratic or cubic terms, interaction terms, et cetera).

Thus optimal policy mixes would feature *equalization of the marginal cost-effectiveness ratios,* $\frac{\partial \Gamma_c}{\partial z_j}/(-\gamma_j)$, across all the activated policy tools. The optimal gap between expected GDP growth and GaR would be decreasing in both the common cost-effectiveness ratio and the aversion to financial instability.

With interactions between tools, the optimal gap may no longer be constant since the compound effectiveness of a given policy mix may depend on the intensity with which policies are activated. For example, if a rise in risk variable x calls for a more intensive use of two complementary policies that jointly exhibit decreasing returns to intensity (akin to when complementary inputs are combined in a production function with decreasing returns to scale), then the optimal policy will accommodate (as in the case of a single policy variable with decreasing marginal effectiveness discussed above) the decreasing effectiveness by tolerating a larger gap when the risk is high than when the risk is low.

4.2.3 Interaction with other policies

The discussion so far has not explicitly dealt with the case in which policies other than macroprudential policies have an impact on expected growth and GaR. If such policies are beyond the control of the macroprudential policy maker (e.g. because of being the responsibility of a separate authority with its own mandate), a way to integrate them into the framework would be to add variables representing those policies in a vector version of risk variable x. Under this reformulation, x would then account not only for risk variables in a narrow sense but also any other factor relevant to the policy design problem of the macroprudential policymaker. In this formulation (which would resemble other setups in which authorities controlling different policy tools interact as in a non-cooperative game), the state of other relevant non-macroprudential policies would enter as part of a vector version of x in the macroprudential policy rule (9). The policy rule could then be interpreted as the macroprudential policy reaction function reflecting the "best response" of macroprudential policy to the settings of other policies.

A more general discussion covering the issue of optimal policy coordination would require further extensions. Discussing all suitable specifications of the problem would exceed the scope of this paper. Eventually, the preferred specification will depend on the capacity to, first, empirically capture the relevant interactions between the policies and their impact on the final policy objectives and, second, effectively coordinate the policies. For instance, to study optimal coordination with monetary policy, the objective function W might have to add terms to reflect goals of this policy (e.g. price stability) not be fully captured by those in (5).[23] A policymaker simultaneously optimizing on the two policies would add the non-macroprudential policy under consideration as an element of a vector version of policy variable z, giving rise to issues similar to those discussed in subsections 4.2.1 and 4.2.2. These extensions might also allow to study the implications of having several authorities with separate objectives acting on each policy in a non-cooperative manner.

4.2.4 Example with interaction between macroprudential and monetary policies

As a simple illustration of the approach, consider the case of policy coordination when the overall policy objective can be represented as in (5) and the vector of relevant policies only includes a macroprudential policy variable z_1 and a monetary policy variable z_2. Then, in the spirit of the specifications explored in Brandao-Marques et al. (2020), the equations for expected growth and GaR could be specified as follows:

$$\bar{y} = \alpha + \beta x + \gamma_1 z_1 + \gamma_2 z_2 + \gamma_3 z_1 z_2, \tag{30}$$

and

$$y_c = \alpha_c + \beta_c x + \gamma_{c1} z_1 + \gamma_{c2} z_2 + \gamma_{c3} z_1 z_2. \tag{31}$$

Assuming the parameters are such that the optimal policy mix is interior, the following FOCs would have to be satisfied:

$$(\gamma_i + \gamma_3 z_{-i}) - w(\bar{y} - y_c)[(\gamma_i + \gamma_3 z_{-i}) - (\gamma_{ci} + \gamma_{c3} z_{-i})] = 0,$$

for policies $i = 1, 2$ and with $-i$ denoting the policy other than i. This implies that the optimal target would relate to the optimal policies as follows:

$$\bar{y} - y_c = \frac{1}{w} \frac{1}{1 + (\gamma_{ci} + \gamma_{c3} z_{-i})/(-\gamma_i - \gamma_3 z_{-i})} \tag{32}$$

for $i = 1, 2$, and also that the two policy tools should co-move respecting the relationship:

$$\frac{\gamma_{c1} + \gamma_{c3} z_2}{-\gamma_1 - \gamma_3 z_2} = \frac{\gamma_{c2} + \gamma_{c3} z_1}{-\gamma_2 - \gamma_3 z_1}, \tag{33}$$

whereby the marginal cost-effectiveness of the two policies remains aligned as, for example, risk variable x moves.[24] In fact (33) implies that the optimal interior values of z_1 and z_2 lie on a line with

$$z_2 = \frac{\gamma_2 \gamma_{c1} - \gamma_1 \gamma_{c2}}{\gamma_{c2} \gamma_3 - \gamma_2 \gamma_{c3}} + \frac{\gamma_{c1} \gamma_3 - \gamma_1 \gamma_{c3}}{\gamma_{c2} \gamma_3 - \gamma_2 \gamma_{c3}} z_1,$$

[23] Examples of this approach include Cecchetti and Kohler (2014), who consider coordination between conventional monetary policy and capital regulation in a related reduced-form setup, and Brandao-Marques et al. (2020), who assess the net gains from historically applied policies using a loss function that adds up the square of deviations of output growth from potential output growth and the square of the inflation rate.

[24] Notice that, other things being equal, a change in x would affect $\bar{y} - y_c$ but not the other terms in (32), thus calling for a re-adjustment in the optimal values of z_1 and z_2.

on which, in principle, dz_2/dz_1 might have any sign, depending on the sign and relative size of the relevant parameters.

In the particular case in which both policies feature $\gamma_i < 0 < \gamma_{ci}$ thus exhibiting, solely considered, the same qualitative policy trade-off as macroprudential policy in the benchmark formulation, having $\gamma_3 \geq 0$ and $\gamma_{c3} \geq 0$, with at least one of these inequalities being strict, would make the two policies complementary (that is, positively co-moving) in response to risk changes. Instead, if $\gamma_3 \leq 0$ and $\gamma_{c3} \leq 0$, with at least one of these inequalities being strict, the two policies would be substitutes (that is, negatively co-moving).

This discussion illustrates the wealth of results that could arise after allowing for interaction between policies. However, in practical applications involving interacting policies, the implications for optimal policy design should be managed with great care. Some parameter values may push the candidate interior solution out of the admissible range of variation of the policy variables or make the candidate solution not satisfy the relevant second order conditions for maximizing W.[25] In these cases, the actual optimal solution would not be interior in that at least one of the policies should remain at either is upper or its lower bound (at least from the point of view of the objectives reflected in W).

Summing up, the GaR approach as described in this paper does not per se presume the existence or inexistence of margins along which macroprudential policy and monetary policy might or might not be complementary to each other. On the contrary, it offers a framework where such complementaries (or the lack of them) could be formalized, estimated, and analyzed. The final question, however, is genuinely empirical. The approach per se cannot establish detailed prescriptions regarding the optimal policy mix without relying on data, specific parameter estimates, and careful analysis of the policy problem under such parameters.

4.3 Policies measured with discrete variables

This subsection discusses the case in which policy variables are discrete. To keep the discussion simple, we articulate it around the benchmark model with a single policy variable z whose effect on expected growth and GaR is as specified in (2) and (3). Intuitively, to characterize the optimal value of the policy when z can take a countable number of values, the left hand side of the FOC in (7) must be replaced by its finite differences counterpart and its sign checked to discover whether there are gains from increasing (or keeping increasing) the variable or, conversely, whether there could be gains from reducing it.

[25] The discussion in footnote 19 applies here too.

More formally, assume that the policy variable can take N different values: $z \in \{z_1, z_2, ...z_N\}$ with $N \geq 2$ and let

$$\Delta W(x, z_i) = W(x, z_{i+1}) - W(x, z_i) \tag{34}$$

represent the welfare gain from increasing the discrete policy variable by one notch when starting from z_i. Using the definition of W in (5) and the expressions for (2) and (3), we obtain the following expression:

$$\Delta W(x, z_i) = \gamma(z_{i+1} - z_i) - \frac{w}{2}(\gamma_c - \gamma)^2(z_{i+1}^2 - z_i^2) + w(\gamma_c - \gamma)A(x)(z_{i+1} - z_i), \tag{35}$$

with $A(x) = (\alpha - \alpha_c) + (\beta - \beta_c)x > 0$. Under the assumptions in (4), the first two terms in this expression are negative, reflecting the direct expected GDP cost of tightening macro-prudential policy and the impact of this cost in reducing the gap between expected growth and GaR, which diminishes the marginal gains from further tightening. The third term is positive and increasing in risk variable x and captures the gap reducing gains from tightening the policy. In a typical case, $\Delta W(x, z_i)$ will be positive at low values of i and turn negative at higher values of i, identifying the optimal policy as the highest i for which $\Delta W(x, z_i)$ is positive. Intuitively, as $A(x)$ is increasing, the optimal level of activation of the discrete policy will generally be higher for higher values of the risk variable x.

A particular case of interest in some applications is that in which the possible values of the policy variable are equally spaced (e.g. when using a cumulative index of macroprudential policy actions). If the scale of the variable is normalized to make the space between any two consecutive values to be one and sets $z_1 = 0$, then $z_i = i - 1$ and we can use $z_{i+1}^2 - z_i^2 = 2i - 1$ to write

$$\Delta W(x, z_i) = \gamma - \frac{w}{2}(\gamma_c - \gamma)^2(2i - 1) + w(\gamma_c - \gamma)A(x), \tag{36}$$

whose negative second term depends linearly on i reflecting, all other things being equal, diminishing marginal welfare gains from the activation of discrete policy at higher and higher levels.

In the even more special case where the policy variable z is binary and can only take values of 0 (inactive) or 1 (active), the welfare gain from activating the policy can be found setting $i = 0$ in (36):

$$\Delta W(x, 0) = \gamma - \frac{w}{2}(\gamma_c - \gamma)^2 + w(\gamma_c - \gamma)A(x), \tag{37}$$

whose interpretation is the same as that provided for the more general case.

In terms of Figure 1, the discreteness of the policy variable does not alter the indifference curves and the location of the "hypothetical" policy frontier that would emerge if z were continuous. The difference is that the effective frontier now only includes as many points on the hypothetical frontier as possible values for z_i. Heuristically, it is still correct to think about the optimal policy as the one that brings the gap between expected growth and GaR as close as possible to the gap in (16) that would be targeted if z were a continuous variable.

5 Further discussion

This section considers variations of the benchmark model that help qualify the importance of the underlying assumptions and the relationship of the setup with alternative approaches to macroprudential policy assessment and design.

5.1 What if the policy variable seems to involve no trade-off?

Let us suppose that an empirical implementation of the quantile regression methodology as represented by (2) and (3) yields an estimate of the impact of policy z on expected growth, γ, not significantly different from zero, while the estimate of the impact of z on GaR, γ_c, is significantly positive. Under the remaining assumptions of the benchmark model, these estimates would imply the inexistence of a policy trade-off. Thus, the policy variable z should be increased up to the point in which either the policy variable reaches its upper bound or the gap between expected growth and GaR becomes zero, whichever happens first.

If the policy variable has a natural upper limit (e.g., it is a binary or discrete variable measuring the activation or deactivation of a buffer measure of fixed size, or qualitative-in-nature policies such as the adoption of effective resolution regimes or other structural reforms), then the implication that the policy should be activated at its maximum level may be meaningful and require no further adjustment in the analysis. In contrast, in the case z represents a policy that can be activated at, in principle, arbitrarily high levels, the possibility that it can lead to fully close the gap between expected growth and GaR does not seem plausible: it would imply that a single tool can fully eliminate all fluctuations in GDP growth over the policy horizon. This suggests some undue extrapolation of estimated effects that, possibly, are only valid in a local sense: as linear approximations to the (average) effects that a more complicated, non-linear relationship implies in the neighborhood of data points used in the estimation. For example, it can happen that within the range of variation of z found in the data, policy does not have on-average detrimental effects on expected growth but that at sufficiently high levels of activation those effects would become visible. Additionally, policies might have never been historically activated a high enough levels to detect their detrimental effects on expected growth.

Practical solutions to these problems may involve running non-linear specifications of (2) or, if the available data does not allow the conjectured non-linearity to be captured, introducing the suspected missing growth cost of the policy at high levels of activation using an auxiliary calculation. This auxiliary calculation might come from structural models or models based on granular data (e.g. like those use in quantitative impact studies of regulatory reforms) that suggest that the measure could have a contractionary impact.

A similar hybrid-modelling approach could be used to introduce into the GaR approach detrimental welfare effects of macroprudential policies that do not necessarily translate into lower expected growth but, for example, into an increase in income or wealth inequality. If consistent with the mandate of the macroprudential authority, the corresponding (otherwise missing) estimated marginal certainty-equivalent welfare cost of the policy (expressed as a fraction of initial GDP) could be added to the equation for expected GDP growth. With this adjustment, the design and assessment of macroprudential policy could proceed as indicated in the previous sections.

5.2 Intermediate objectives and targeted policy tools

Current practice of macroprudential policy largely involves a disaggregated approach. Authorities around the globe, as well as research in the field of macroprudential policy, often address the design and assessment of macroprudential policy by splitting them into separate dimensions.[26] As in microprudential regulations, these dimensions are commonly determined by the nature of the underlying source of systemic risk (e.g. liquidity vs. solvency risk) or by the sector that originates, transmits or suffers the risk (e.g. banks vs. non-banks, commercial vs. residential real estate, etc.). The resulting silos are typically associated with an intermediate objective and one or several dedicated policy tools (e.g. "capital-based tools for the banking sector", "liquidity-based tools for the investment management sector" or "borrower-based tools for residential real estate risk").

The practical attractiveness of the disaggregated approach stems from the difficulties of integrating under a common general equilibrium perspective and with a common ultimate goal the multiple aspects of systemic risk, the various factors contributing to financial stability (or the lack of it), and the rich sets of policy tools available to address these dimensions

[26] The defense of the disaggregated approach is mostly found in policy references. See, for instance, IMF (2013) which defends the convenience of having well-defined (intermediate) objectives (and tools associated with each of them) at the time of operationalizing macroprudential policies (e.g., on pages 19-20 and 29). More recently, the Financial Stability Board (2021) has argued in favor of articulating the periodic assessment of systemic risk by macroprudential authorities around a number of relevant vulnerabilities. This perspective is consistent with abundant research work identifying each of the vulnerabilities as drivers or amplifiers of systemic financial crises.

and factors. The purpose of this subsection is to show that the disaggregated approach is not incompatible with the analytical framework and empirical efforts associated with the GaR approach. In fact, the latter can contribute to integrating, adding up or at least putting under a common umbrella sectoral macroprudential policies that might, otherwise, be difficult to relate to each other when trying to obtain an overall notion of macroprudential policy stance.

To illustrate the compatibility between the GaR approach and the disaggregated approach, we consider a variation of our benchmark model in which macroprudential policy involves M different dimensions, $j = 1, 2, ..M$. The variation presented here relies on simplifying assumptions that are directed at making the formulation analytically tractable rather than empirically plausible. In this spirit, we assume that each dimension j can be associated with an intermediate objective I_j and a targeted policy tool z_j.[27] We further assume that intermediate objectives can be represented as linear functions of their targeted tools

$$I_j = \lambda_{0j} + \lambda_{1j} z_j, \tag{38}$$

where λ_{0j} is an autonomous component of the intermediate objective and $\lambda_{1j} > 0$ measures the marginal impact of the targeted policy variable on the intermediate objective.

The baseline equations (2) and (3) could then be reformulated as follows:

$$\bar{y} = \alpha + \Sigma_{j=1}^{M} \gamma_j z_j, \tag{39}$$

and

$$y_c = \alpha_c + \Gamma_c(I_1, I_2, ...I_M), \tag{40}$$

where $\gamma_j < 0$ for all j and Γ_c is an increasing and strictly concave function of the vector of intermediate objectives.[28] Therefore, as in the baseline setup, macroprudential policies involve a trade-off: increasing policy z_j improves the intermediate objective I_j but at the cost of reducing mean growth at the margin.

[27] This formulation should not be regarded as a proposal to undertake the GaR approach with explicit reference to intermediate objectives. That would require identifying the causal effects of each tool on the intermediate targets and of each (endogenous, simultaneously determined) intermediate target on output growth. This can be econometrically more challenging than directly inferring the (reduced form) causal effects of policy actions (or policy shocks) on output growth.

[28] To easen the presentation, we assume monotonicity in the impact of each I_j on $\bar{y}$ and y_c, and impose a sign convention for I_j and z_j such that increasing I_j is *good* for financial stability (that is, increases y_c) and increasing the policy variable z_j is *good* for the intermediate objective j (that is, increases I_j). Additionally we do not explicitly include any risk variable x in the equations. However, it would be trivial to introduce one or a vector of them affecting $\bar{y}$ and y_c linearly as in the baseline model. We could also consider risk variables that affect the autonomous component λ_{0j} of each intermediate objective.

As in the non-linear world described in the extension with optimal policy mixes, all targeted policy variables activated at a strictly positive level at the optimum would satisfy a modified version of (7) and, consequently, (16), implying

$$\bar{y} - y_c = -\frac{1}{w} \frac{\frac{\partial \bar{y}}{\partial z_j}}{\frac{\partial \bar{y}}{\partial z_j} - \frac{\partial y_c}{\partial z_j}} = \frac{1}{w} \frac{1}{1 + \frac{\partial \Gamma_c}{\partial I_j} \lambda_j / (-\gamma_j)}. \tag{41}$$

Thus, the optimal vector of targeted policies would again feature equalization of the marginal growth-gap rates of transformation and, hence, *equalization of the marginal cost-effectiveness ratios* across all activated policy tools. Moreover, the optimal gap between expected GDP growth and GaR would be decreasing in both the aversion to financial instability and the common ratio. The cost-effectiveness ratio of targeted policy j is the ratio between the marginal effectiveness of the policy, that is, its marginal capability to improve GaR by affecting intermediate objective j ($\frac{\partial \Gamma_c}{\partial I_j} \lambda_j$) and the marginal cost of the policy in terms of mean growth ($-\gamma_j$).

Depending on the degree to which intermediate objectives may feature complementarity in their compound impact on GaR, as captured by the cross-derivatives of function Γ_c, the setup with multiple intermediate objectives might imply increasing or decreasing the target gap as well as varying the optimal policy mix in response to changes in, for instance, the autonomous component of one of the intermediate objectives. For example, in the simple case in which Γ_c is additively separable across intermediate objectives (so that they do not directly interact in affecting GaR), the policy response to a deterioration in the autonomous component of one objective j would be the tightening of policy across all intermediate objectives (so that $\frac{\partial \Gamma_c}{\partial I_{j'}}$ declines across all policy dimensions j' and the equality in (41) is restored at a higher target gap).

5.3 GaR vs. directly focusing on systemic financial crisis

Macroprudential policy is commonly associated with the ultimate goal of minimizing the frequency and severity of systemic financial crises. The specificity of this objective helps differentiate this policy from other policies such as those pursuing price stability or the dampening of business cycle fluctuations. Assessing and designing macroprudential policy under the GaR approach builds on the empirically documented negative impact of systemic financial distress (and the previous build-up of systemic financial vulnerabilities) on (future) growth outcomes, and especially on the lower tail of the distribution of those outcomes. However, neither the approach nor the supportive evidence imply that the only reason why the economy may experience very low growth outcomes are financial crises or the materialization of systemic risk.

The purpose of this section is to explore the connection between the benchmark model with which we have illustrated the GaR approach and a conceptual framework which explicitly distinguishes "normal times" from "systemic crises" while allowing for the possibility of reaching both good and bad growth outcomes in each state. We aim to show that a framework in which macroprudential policy explicitly reduces the probability and severity of systemic crises is compatible with the reduced-form formulation of the GaR approach.

To see this, consider the situation in which growth outcomes y over the policy horizon follow a two-regime process specified as follows:

$$y = \begin{cases} y_0 - g_H(z), & \text{with prob. } 1 - \varepsilon(z), \\ y_0 - \Delta + g_L(z), & \text{with prob. } \varepsilon(z), \end{cases} \tag{42}$$

where y_0 is a random variable representing some baseline stochastic growth rate, z represents macroprudential policy, $\varepsilon(z)$ is the probability of a systemic financial crisis, Δ is the baseline growth cost of a systemic crisis, $g_H(z)$ is a deterministic function that measures the growth cost of macroprudential policy in "normal times" (a regime identified by the subscript H), and $g_L(z)$ is a deterministic function that measures the growth benefit of macroprudential policy during "systemic crises" (a regime identified by the subscript L). Accordingly, growth in normal times is $y_H(z) = y_0 - g_H(z)$, while growth in a systemic crisis is $y_L(z) = y_0 - \Delta + g_L(z)$. We assume $\Delta > g_H(z) + g_L(z)$ so as to have $y_H(z) > y_L(z)$ for all relevant z.

To set the basis for a policy trade-off in this setting, we assume that tightening the policy z can reduce the probability of a systemic crisis ($\varepsilon' < 0$) and and its severity ($g_L' > 0$) but only at a cost (at least beyond some point) in terms of growth in normal times ($g_H' > 0$). Finally, we assume that y_0 has a density function $f(\cdot)$ and a cumulative distribution function $F(\cdot)$.

In this setup, expected growth is

$$\bar{y} = E(y_0) - (1 - \varepsilon(z))g_H(z) - \varepsilon(z)(\Delta - g_L(z)), \tag{43}$$

which declines with z provided that

$$(1 - \varepsilon(z))g_H'(z) - \varepsilon(z)g_L'(z) > -\varepsilon'(z)(\Delta - g_L(z) - g_H(z)), \tag{44}$$

where the expression in the right hand side is positive. In other words, having $d\bar{y}/dz < 0$ as in the benchmark GaR model of prior sections requires policy to have a cost in terms of expected growth (for a given crisis probability $\varepsilon(z)$) larger than the expected growth gains implied by the reduction in the probability of a crisis. Other things being equal, this condition is more likely to hold the larger is the cost of the policy in terms of normal-times growth, $g'_H(z)$, and the lower is the probability of a crisis, $\varepsilon(z)$.

The GaR in this setup can in turn be found as the relevant c quantile of the growth rate, that is, the growth rate y_c that solves

$$\Pr(y \leq y_c) \equiv (1-\varepsilon(z))\Pr(y_H(z) \leq y_c) + \varepsilon(z)\Pr(y_L(z) \leq y_c) = c, \tag{45}$$

where the first identity follows from the Law of Total Probability. Using previous definitions, the equation that implicitly defines y_c can be written as

$$(1-\varepsilon(z))F\left(y_c + g_H(z)\right) + \varepsilon(z)F\left(y_c + \Delta - g_L(z)\right) = c. \tag{46}$$

Using the Implicit Function Theorem, one can totally differentiate (46) to find the marginal impact of policy z on GaR, dy_c/dz. It is immediate to check that $dy_c/dz > 0$, as in the benchmark GaR model of prior sections, if and only if

$$\begin{aligned} -\varepsilon'(z)\left[F(y_c + \Delta - g_L(z)) - F(y_c + g_H(z))\right] \;>\; & (1-\varepsilon(z))f(y_c + g_H(z))g'_H(z) \\ & -\varepsilon(z)f(y_c + \Delta - g_L(z))g'_L(z), \end{aligned} \tag{47}$$

where the left hand side is positive because $\varepsilon' < 0$, $\Delta > g_H(z) + g_L(z)$, and $F(\cdot)$ is an increasing function. The right hand side of (47) can in principle be positive or negative depending on the relative size of each of its two terms. If it is negative, then (47) definitely holds and $dy_c/dz > 0$ as in the benchmark GaR setup of prior sections. Otherwise, satisfying (47) requires that its right hand side is relatively small, which is more likely to be the case (i) if $g'_H(z)$ is small relative to $g'_L(z)$ (macroprudential policy is sufficiently cost-effective) and (ii), provided that the density function $f(\cdot)$ is increasing at its low tail, if Δ large relative to $g_H(z) + g_L(z)$ (financial crises have a severe impact on growth outcomes).

Summing up, under plausible conditions, a semi-structural model in which GDP growth is explicitly affected by systemic financial crises and macroprudential policy helps reduce the probability and severity of such crises has implications for $d\bar{y}/dz$ and dy_c/dz similar to those postulated in the reduced-form GaR setup of prior sections. Thus the equations of the empirical GaR formulation can be interpreted as local approximations to the relationship linking the policy variable z to expected growth and GaR in the semi-structural model. The semi-structural model would, in that sense, provide a structural foundation or explanation for why containing the incidence and severity of systemic risk using macroprudential policy is positive for GaR but may, perhaps beyond some level of activation, reduce the unconditional expected growth rate. Yet, as in other empirical problems, the advantages of using a reduced-form approach instead of a more structural approach may ultimately depend on data availability and the extent to which the modeler knows the "correct" structural model.[29]

5.4 Accounting for the term structure of macroprudential policies

Consider an infinite horizon dynamic version of the benchmark model in which the policy-maker at any period t aims to maximize the expected present discounted value of welfare flows of all subsequent periods, that is,

$$W_t = E_t \left\{ \sum_{s=1}^{\infty} \Lambda^s \left[\bar{y}_{t+s} - \frac{1}{2} w (\bar{y}_{t+s} - y_{c,t+s})^2 \right] \right\}, \tag{48}$$

where Λ is the discount factor and the welfare flows in each future period $t+s$ are determined exactly like the single-shot welfare W of the benchmark model: by the expected growth rate and the (expected) gap between the expected growth and the GaR in the corresponding period. Consider further a generalized version of (2) and (3) in which the policy adopted in each period t has the potential to affect expected growth and GaR in each subsequent period, that is,

$$\bar{y}_{t+s} = \alpha + \beta x_{t+s-1} + \sum_{l=1}^{\infty} \gamma_l z_{t+s-l}, \tag{49}$$

and

$$y_{c,t+s} = \alpha_c + \beta_c x_{t+s-1} + \sum_{l=1}^{\infty} \gamma_{c,l} z_{t-s-l}. \tag{50}$$

Under this formulation, risk variable x_{t+s-1} can be interpreted as summarizing the one-period ahead predictors of GDP growth at $t+s$ other than the policies z_{t+s-l} adopted at $t+s-l$ for all lags $l = 1, 2, ...$ (that is, in all periods prior to $t+s$). In turn, γ_l and $\gamma_{c,l}$ measure the impact of the policy adopted l periods before, z_{t+s-l}, on the expected growth and GaR at each period $t+s$.

So, when solving the optimal policy problem at t, the policy maker takes into account the impact of the policy adopted at t on each of welfare flows of the $t+s$ periods ahead over which such policy will have an impact on growth outcomes (that is, for which γ_s or $\gamma_{c,s}$ are different from zero). The FOC for an optimal interior choice of z_t in this problem, $\partial W_t / \partial z_t = 0$, can be written as

$$\sum_{s=1}^{\infty} \Lambda^s \gamma_s - w \sum_{s=1}^{\infty} \Lambda^s (\gamma_s - \gamma_{c,s}) E_t (\bar{y}_{t+s} - y_{c,t+s}) = 0. \tag{51}$$

[29] In the case at hand, a more structural approach would require getting reliable estimates of the function $\varepsilon(z)$ which describes the impact of policy z on the probability of a systemic financial crisis, as well as the functions $g_H(z)$ and $g_L(z)$ that describe the costs and benefits (in terms of growth) of the policy in normal and crises times, respectively.

But, then, there exists a constant target gap defined as

$$(\bar{y} - y_c)^* = \frac{1}{w} \frac{\sum_{s=1}^{\infty} \Lambda^s \gamma_s}{\sum_{s=1}^{\infty} \Lambda^s (\gamma_{c,s} - \gamma_s)} = \frac{1}{w} \frac{1}{1 + \left(\sum_{s=1}^{\infty} \Lambda^s \gamma_{c,s}\right) / \left[\sum_{s=1}^{\infty} \Lambda^s (-\gamma_s)\right]} \tag{52}$$

such that:

1. A policy rule setting

$$z_{t+s-1} = \frac{E_{t+s-1}(\bar{y}_{t+s} - y_{c,t+s} \mid z_{t+s-1} = 0) - (\bar{y} - y_c)^*}{\gamma_{c,1} - \gamma_1} \tag{53}$$

   for all $s = 1, 2, ...$, allows to induce $E_{t+s-1}(\bar{y}_{t+s} - y_{c,t+s}) = (\bar{y} - y_c)^*$ at every period $t + s - 1$ irrespectively of the initial conditions at $t + s - 1$.

2. The path of $E_t(\bar{y}_{t+s} - y_{c,t+s})$ induced by the policy rule (53) satisfies $E_t(\bar{y}_{t+s} - y_{c,t+s}) = (\bar{y} - y_c)^*$ for all s and, thus, the FOC in (51) for all t.

Therefore, the policy rule in (53) is optimal and induces a constant gap $(\bar{y} - y_c)^*$ between expected growth and GaR in every future period.[30]

Importantly, the target gap $(\bar{y} - y_c)^*$ as defined in (52) is qualitatively determined by the same trade-offs as in the one-shot benchmark model: it is inversely proportional to the policymaker's aversion to financial instability w and decreasing in the marginal intertemporal cost-effectiveness of the policy tool, that is, the ratio of the discounted intertemporal capacity of the policy tool to improve GaR in subsequent periods, $\sum_{s=1}^{\infty} \Lambda^s \gamma_{c,s}$, to the discounted intertemporal cost in terms of expected growth, $\sum_{s=1}^{\infty} \Lambda^s (-\gamma_s)$.

This extension suggests that the short-cut taken in the one-shot benchmark formulation, that is, considering just the (marginal) effects γ and γ_c of policy on the cumulative-over-the-policy-horizon expected growth and GaR, respectively, is a good approximation to what a full intertemporal formulation would imply provided that (i) most of the effects of the adopted policy occur within the policy horizon (rather than with a lag longer than the length of the policy horizon) and (ii) the discount factor Λ is close to 1.

6 Conclusions

Using the concept of GaR in the measurement of the downside risks that macroprudential policy aims to address opens very interesting avenues for the use of empirical quantitative

[30] The property of optimally reaching the target $(\bar{y} - y_c)^*$ in every period would be replaced by a process of gradual dynamic approximation to the target in the presence of costs to the adjustment of the policy variable between periods. Similar richer dynamics might also arise in the presence of other non-linear interactions between policies adopted in different periods or in cases, such those discussed at the end of subsection 4.2.3 in which the welfare criterion is expanded to include other objectives.

models in the design of macroprudential policies and the development of concrete notions of macroprudential policy stance. The setup allows us to explicitly consider, using a similar modeling methodology, the effects of risk and policy variables on expected GDP growth (arguably, a succinct measure of what other macroeconomic policies care about) and the risk of sufficiently adverse GDP growth outcomes (arguably, a promising concrete measure of what macroprudential policy cares about). This paper has explored the foundations for the design and assessment of macroprudential policies using this setup.

The paper starts with a stylized benchmark description of the setup in the context of its implementation using the outcome of quantile regressions. A welfare criterion for the design of the optimal policies has been proposed that can be microfounded as consistent with the maximization of the expected utility of a representative agent in some contexts. The properties of the optimal policies have been explored in the benchmark setup as well as in several extensions that bring the setup closer to what might be relevant in empirical applications and/or show the robustness and compatibility of the approach with alternative approaches. The extensions cover cases with non-linearities in the impact of policy variables and risk variables on the relevant outcomes, interactions between multiple policy tools, and discrete policy variables. Additional discussions consider the case in which policies that seem to involve no trade-off between mean growth and GaR, the compatibility of the framework with the view that macroprudential policy involves several intermediate objectives and policy tools targeted to such objectives, and explore the connection, the relationship with a framework focused on preventing and mitigating the effects of systemic financial crises, and the explicit inclusion of intertemporal effects in the policy design problem.

Under the postulated representation of preferences, the policy design problem yields a quantitative-based policy target and a metric for the assessment of policy stance similar to that of other macroeconomic policies. While empirically-relevant specifications of the problem will likely be more complicated than the benchmark model explored in this paper and, thus, require adapting the analysis accordingly, the main challenges for the applicability of the explored framework are more empirical and political than conceptual. On the empirical side, the main challenge resides in the consistent and sufficiently precise estimation of the causal effects of risk and policy variables on the relevant moments (mean and GaR) of the growth distribution. Properly detecting relevant non-linearities and interactions between policies is also important. In the absence of proper estimates of the relevant parameters and relationships, the mechanical application of this framework could produce misguided policy advice. Therefore, the framework will develop at the speed with which data on the applied policies accumulates and econometric efforts succeed in providing reliable estimates of their effects on growth outcomes.

On the political side, once data and estimation provide a reliable description of the policy trade-offs, the main challenge would be to define society's aversion for financial instability on which optimal policies should be based. Additionally, given the uncertainty surrounding the relevant parameters implied by the empirical challenges, policymakers may need to be guided on how to expand the type of framework sketched in this paper to account for model uncertainty (that is, for the imperfect knowledge of the specification and parameters of the relevant quantile regressions) and the potential policy mistakes that could stem from this uncertainty.

References

Adrian, T., N. Boyarchenko, and D. Giannone (2019) "Vulnerable growth," *American Economic Review*, 109, 1263-89.

Adrian, T., F. Grinberg, N. Liang, and S. Malik (2018) "The term structure of Growth-at-Risk," International Monetary Fund Working Paper 18/180.

Aikman, D., J. Bridges, S. Burgess, R. Galletly, I. Levina, C. O'Neill, and A. Varadi (2018) "Measuring risks to UK financial stability," Bank of England Staff Working Paper 738.

Aikman, D., J. Bridges, S. Hacioglu, C. O'Neill, and A. Raja (2021) "Credit, capital and crises: a GDP-at-Risk approach," CEPR DP 15864.

Arbatli-Saxegaard, E., K. Gerdrup, and R. Johansen (2020) "Financial imbalances and medium term growth-at-risk in Norway," Norges Bank Staff Memo No. 5-2020.

Brandao-Marques, L., G. Gelos, M. Narita, and E. Nier (2020) "Leaning against the wind: A cost-benefit analysis for an integrated policy framework," International Monetary Fund Working Paper 20/123.

Brownlees, C., and A. Souza (2021) "Backtesting global Growth-at-Risk," *Journal of Monetary Economics*, 118, 312-330.

Caldera-Sánchez, A., and O. Röhn (2016) "How do policies influence GDP tail risks?," OECD Economics Department Working Papers 1339.

Cecchetti, S. (2008) "Measuring the macroeconomic risks posed by asset price booms," in J. Campbell, editor, *Asset Prices and Monetary Policy*, University of Chicago Press for NBER.

Cecchetti, S., and M. Kohler (2014) "When capital adequacy and interest rate policy are substitutes (and when they are not)," *International Journal of Central Banking*, 10, 205-232.

Cecchetti, S., and H. Li (2008) "Measuring the impact of asset price booms using quantile vector autoregressions," Brandeis University, mimeo.

Cecchetti, S., and J. Suarez (2021) "On the stance of macroprudential policy," Reports of the Advisory Scientific Committee No. 11, European Systemic Risk Board, December.

Chavleishvili, S., R. Engle, S. Fahr, M. Kremer, S. Manganelli, and B. Schwaab (2020) "The risk management approach to macro-prudential policy," European Central Bank, mimeo.

De Nicolo, G., and M. Lucchetta (2017) "Forecasting tail risks," *Journal of Applied Econometrics*, 32, 159-170.

Duprey, T., and A. Ueberfeldt (2020) "Managing GDP Tail Risk," Bank of Canada Staff Working Paper 2020-03.

Elenev, V., T. Landvoigt, and S. V. Nieuwerburgh (2021) "A macroeconomic model with financially constrained producers and intermediaries," *Econometrica*, 89, 1361-1418.

ESRB (2021), A Framework for Assessing Macroprudential Stance, Report of the Expert Group on Macroprudential Stance – Phase II (implementation), European Systemic Risk Board, December.

Figueres, J. M., and M. Jarocinski (2020) "Vulnerable growth in the euro area: Measuring the financial conditions," *Economics Letters*, 191, 109126.

Financial Stability Board (2021), FSB Financial Stability Surveillance Framework, Financial Stability Board, September.

Fishburn, P. (1977) "Mean-risk analysis with risk associated with below-target returns," *American Economic Review*, 67, 116-126.

Franta, M., and L. Gambacorta (2020) "On the effects of macroprudential policies on Growth-at-Risk," *Economics Letters*, 196, 109501.

Gadea Rivas, M. D., L. Laeven, and G. Perez-Quiros (2020) "Growth-and-risk trade-off," ECB Working Paper No. 2397.

Galán, J. (in press) "The benefits are at the tail: uncovering the impact of macroprudential policy on GDP growth," *Journal of Financial Stability*, in press. https://doi.org/10.1016/j.jfs.2020.100831.

Galati, G., and R. Moessner (2018) "What do we know about the effects of macroprudential policy?" *Economica*, 85, 735-770.

International Monetary Fund (2013), "Key aspects of macroprudential policy," International Monetary Fund Policy Paper, June.

Kahneman, D., and A. Tversky (1979) "Prospect theory: An analysis of decision under risk," *Econometrica*, 47, 263-291.

Kilian, L., and S. Manganelli (2008) "The central banker as a risk manager: Estimating the Federal Reserve's preferences under Greenspan," *Journal of Money, Credit and Banking*, 40, 1103-1129.

Laeven, L., and F. Valencia (2018) "Systemic banking crises revisited," International Monetary Fund Working Paper 18/206.

MAG (2010), Assessing the Macroeconomic Impact of the Transition to Stronger Capital and Liquidity Requirements – Final Report, Macroeconomic Assessment Group, Basel Committee on Banking Supervision, December

Mendicino, C., K. Nikolov, J. Suarez, and D. Supera (2018) "Optimal dynamic capital requirements," *Journal of Money, Credit, and Banking*, 50, 1271-1297.

Plagborg-Møller, M., L. Reichlin, and G. Ricco, and T. Hasenzagl (2020) "When is growth at risk?," *Brookings Papers on Economic Activity*, Spring, 167-229.

Poloz, S. (2014) "Integrating uncertainty and monetary policy-making: A practitioner's perspective," Bank of Canada Discussion Paper 2014-6.

Prasad, A., S. Elekdag, P. Jeasakul, R. Lafarguette, A. Alter, A. X. Feng, and C. Wang (2019) "Growth at Risk: Concept and application in IMF country surveillance," International Monetary Fund Working Paper 19/36.

Svensson, L. (2003) "Optimal policy with low-probability extreme events," NBER Working Paper Series No. 10196.

Appendix

A.1 Microfoundations of the GaR-based welfare criterion

Let Y denote GDP and let y describe the implied (geometric) GDP growth rate relative to a benchmark level Y_0 so that

$$Y = (1 + y)Y_0. \tag{54}$$

Suppose also that there is a representative agent whose preferences for GDP levels are represented by a utility function $U(Y)$ with a local coefficient of *absolute* risk aversion $\lambda(Y_0)$ at $Y = Y_0$ and that the utility function can be (locally) described as one exhibiting CARA with parameter $\lambda(Y_0)$, so that

$$U(Y) = -\exp(-\lambda(Y_0)Y). \tag{55}$$

Using (54), we can write

$$U(Y) = -\exp(-\lambda(Y_0)Y_0(1 + y)) = -\exp(-\lambda(Y_0)Y_0)\exp(-\lambda(Y_0)Y_0 y). \tag{56}$$

For fixed Y_0, since affine monotonic transformations of a utility function will represent exactly the same preferences, we can replace $U(Y)$ with

$$u(y) = -\exp(-\lambda(Y_0)Y_0 y) = -\exp(-\rho_0 y), \tag{57}$$

where $\rho_0 = \lambda(Y_0)Y_0$ describes the agent's coefficient of *relative* risk aversion at Y_0. Thus, this utility function describes CARA preferences directly on the growth rate y but the parameter ρ_0 in this specification measures the relative risk aversion of the agent (in terms of their preferences for GDP levels) at the initial GDP level Y_0.

Let us now suppose that GDP growth is normally distributed, so $y \sim N(\bar{y}; \sigma_y^2)$. From the well-known properties of normal distributions, the moment generating function of the distribution of y is then

$$M(t) = E(\exp(ty)) = \exp(\bar{y}t + \frac{1}{2}\sigma_y^2 t^2) \tag{58}$$

for any t. In particular,

$$M(-\rho_0) = E(\exp(-\rho_0 y)) = \exp(-\rho_0\bar{y} + \frac{1}{2}\rho_0^2\sigma_y^2). \tag{59}$$

Hence, from (57) and (59), we can write the agent's expected utility as

$$E[u(y)] = -E[\exp(-\rho_0 y)] = -\exp(-\rho_0\bar{y} + \frac{1}{2}\rho_0^2\sigma_y^2). \tag{60}$$

Further, since monotonic transformations of expected utility will represent exactly the same preferences, these preferences can be equivalently described by the (indirect) utility function

$$v = \bar{y} - \frac{\rho_0}{2}\sigma_y^2. \tag{61}$$

that is, a simple linear expression in the mean $\bar{y}$ and the variance σ_y^2 of the growth rate y.

Growth-at-risk (GaR) for a given confidence level c is the c-quantile of the probability distribution of y, that is, the value y_c such that

$$\Pr(y \leq y_c) = c. \tag{62}$$

Under the properties of normal distributions, $(y - \bar{y})/\sigma_y$ is a standard normal random variable, $N(0, 1)$. If $\Phi(\cdot)$ denotes the cumulative distribution function of a standard normal, we can write

$$\Pr(y \leq y_c) = c \Leftrightarrow \Pr((y - \bar{y})/\sigma_y \leq (y_c - \bar{y})/\sigma_y) = c \Leftrightarrow \Phi((y_c - \bar{y})/\sigma_y) = c. \tag{63}$$

Solving for y_c in the last expression yields

$$y_c = \bar{y} + \sigma_y \Phi^{-1}(c). \tag{64}$$

Alternatively, solving for σ_y yields

$$\sigma_y = \frac{y_c - \bar{y}}{\Phi^{-1}(c)}, \tag{65}$$

which plugged into (61) leads to the indirect utility function

$$v(\bar{y}, y_c; \rho_0, c) = \bar{y} - \frac{\rho_0}{2(\Phi^{-1}(c))^2}(\bar{y} - y_c)^2, \tag{66}$$

which expresses the agent's expected utility as a function of expected growth, GaR at a confidence level c, the relative risk aversion coefficient of the agent at the initial level of GDP ρ_0, and the confidence level c.

Hence, maximizing a welfare criterion of the form

$$W = \bar{y} - \frac{w}{2}(\bar{y} - y_c)^2, \tag{67}$$

as assumed in the main text, would be equivalent to the maximization of the expected utility of the representative agent for

$$w = \frac{\rho_0}{(\Phi^{-1}(c))^2}. \tag{68}$$

For instance, for $c = 0.05$, we have $\Phi^{-1}(c) = -1.6449$, so with a coefficient $\rho_0 = 2$ of relative risk aversion at Y_0, both criteria would coincide under $w = 2(1.6449)^{-2} = 0.7392$.

Intuitively, the policy maker's preference for financial stability should increase with the agent's relative risk aversion parameter ρ_0 as well as, for any $c < 0.5$, with the level of confidence c at which GaR is calculated.[31]

A.2 Modeling GaR vs. growth volatility and departing from normality

Under the normality assumption sustaining the interpretation of the welfare criterion W as consistent with expected utility maximization, modelling a lower quantile such as y_c and expected growth $\bar{y}$ is no different from modeling the standard deviation and the mean of the growth rate and focusing on a welfare criterion that directly depends on those moments of the growth distribution.

Moreover, under normality, if expected growth is determined as in (2) and the standard deviation of the growth rate is linear in x and z, say

$$\sigma_y = \alpha_\sigma + \beta_\sigma x + \gamma_\sigma z, \tag{69}$$

then (64) implies that (69) is exactly compatible with the specification of y_c in (3) if and only if $\alpha_c = \alpha + \Phi^{-1}(c)\alpha_\sigma$, $\beta_c = \beta + \Phi^{-1}(c)\beta_\sigma$, and $\gamma_c = \gamma + \Phi^{-1}(c)\gamma_\sigma$, where for $c < 0.5$ we have $\Phi^{-1}(c) < 0$. So the prior assumption that the policy variable has a positive effect on y_c ($\gamma_c > 0$) and a negative effect on $\bar{y}$ ($\gamma < 0$) would require that it also has a sufficiently large negative impact on σ_y ($\gamma_\sigma < -\gamma/\Phi^{-1}(c) < 0$).

While the ability to structurally interpret the analysis in the main text under the assumption of normality as exactly compatible with expected utility maximization is reassuring, the normal case would not justify a strict preference for the quantile regression approach. A quantile regression approach to the analysis of macroprudential policies is typically defended on the grounds that there are variables whose impact on extreme low quantiles of the growth distribution is empirically detectable, while its impact on the standard deviation of the growth rate (or on high quantiles of the growth distribution) might not be (or at least not so clearly). For instance, it is likely that empirical measures of GDP volatility are dominated by what happens at business cycle frequencies, while what happens at a sufficiently low growth quantile may better capture the impact of infrequent financial crises.

However, representing the world in which lower quantiles are disproportionately affected by one variable or infrequent discrete events have non-linear implications for growth implies departing from the normality assumption and, hence, from the setup in which the interpretation of the welfare criterion in expected utility terms is exactly valid. In other words, while the normal world provides a useful benchmark to help connect the preference for financial stability reflected into the welfare criterion W with a standard way of representing

[31] Notice that, for $c < 0.5$, $\Phi^{-1}(c)$ is negative and approaches zero as c increases, so $(\Phi^{-1}(c))^2$ is decreasing in c.

agents' preferences in economics, it is probably not the most practically relevant one. Box D in Cecchetti and Suarez (2021) describes a number of simulation exercises in which GDP growth is drawn from empirically-motivated non-normal distributions and examines the accuracy with which a GaR-based criterion such as W approximates the true expected utility (measured in certainty equivalent terms). The message from those simulations is that the GaR-based metrics provides a reasonably good approximation to the expected-utility-based metrics even when the growth distribution deviates substantially from normality, as well as under constant relative risk aversion (CRRA) preferences.

Additionally, as discussed in the main text, in the non-normal world, we could interpret W as a heuristic representation of the preferences of a policy maker who cares about the gap $\bar{y} - y_c$, for a suitably low value of c, rather than, for instance, the standard deviation of GDP growth, because of some form of loss aversion. Under this perspective, the focus on the trade-off between maximizing $\bar{y}$ and minimizing the gap $\bar{y} - y_c$ could reflect that the policy maker cares more about the relative output losses incurred at the low tail of the growth rate distribution than the potentially offsetting (in expected terms) relative output gains obtained at the high tail.

A.3 Reformulation using a growth-given-stress criterion[32]

Let us define the growth-given-stress (GgS) for a given reference probability c as the expected value of the GDP growth rate y conditional on this rate being lower than the c-quantile of its distribution y_c, that is

$$GgS_c = E(y \mid y \leq y_c). \tag{70}$$

When y is a normal random variable, GgS_c is just the mean of an upper-truncated normal random variable with truncation point at y_c. The well-known expression for this mean implies

$$GgS_c = \bar{y} - \frac{\phi\left(\frac{y_c - \bar{y}}{\sigma_y}\right)}{\Phi\left(\frac{y_c - \bar{y}}{\sigma_y}\right)}\sigma_y, \tag{71}$$

where $\phi(\cdot)$ is the density function of standard normal. But, since y_c is the c-quantile of the distribution of y, the term $(y_c - \bar{y})/\sigma_y$ can be written as $\Phi^{-1}(c)$. This allows us to write

$$GgS_c = \bar{y} - \frac{\phi\left(\Phi^{-1}(c)\right)}{c}\sigma_y. \tag{72}$$

Using (65) to substitute for σ_y and re-arranging, we can express

$$\bar{y} - GgS_c = \frac{-\phi\left(\Phi^{-1}(c)\right)}{c\Phi^{-1}(c)}\left(\bar{y} - y_c\right), \tag{73}$$

[32] I wish to thank Steve Cecchetti for making me aware of the possibility of this reformulation.

where, for a given c, the ratio $-\phi\left(\Phi^{-1}(c)\right)/(c\Phi^{-1}(c))$ is a proportionality constant (which is positive for $c < 0.5$).

In other words, when the growth rate y is normally distributed the gap between expected growth and GgS is proportional to the gap between expected growth and GaR. Therefore, maximizing the welfare criterion W specified in (5) would be equivalent to maximizing a similar linear-quadratic criterion whose quadratic term contains the square of the distance between expected growth and GgS and where the instability aversion parameter w is replaced by

$$w_{GgS} = \left(\frac{c\Phi^{-1}(c)}{\phi\left(\Phi^{-1}(c)\right)}\right)^2 w. \tag{74}$$

This criterion would thus have the same microfoundation as the criterion provided in Section A1 of this Appendix for W. With this microfoundation, the parameter w_{GgS} would become, using (68),

$$w_{GgS} = \frac{c^2\rho_0}{\phi(\Phi^{-1}(c))} \tag{75}$$

The optimal policy rule resulting from solving the baseline policy problem under the GaR-based welfare criterion would be equivalent to the one that emerges under the equivalent GgS-based criterion, and would also satisfy the constant target gap property in (16). This property could be translated into targeting a gap between expected growth and GgS given by

$$\bar{y} - GgS_c = \frac{1}{w_{GgS}} \frac{1}{1 + \gamma_c/(-\gamma)}. \tag{76}$$